

Qualitative Analysis of VASS-Induced MDPs^{*}

Parosh Aziz Abdulla^{1**}, Radu Ciobanu², Richard Mayr^{2***}, Arnaud Sangnier³, and Jeremy Sproston^{4†}

¹ Uppsala University, Sweden

² University of Edinburgh, UK

³ LIAFA, Univ Paris Diderot, Sorbonne Paris Cité, CNRS, France

⁴ University of Turin, Italy

Abstract. We consider infinite-state Markov decision processes (MDPs) that are induced by extensions of vector addition systems with states (VASS). Verification conditions for these MDPs are described by reachability and Büchi objectives w.r.t. given sets of control-states. We study the decidability of some qualitative versions of these objectives, i.e., the decidability of whether such objectives can be achieved surely, almost-surely, or limit-surely. While most such problems are undecidable in general, some are decidable for large subclasses in which either only the controller or only the random environment can change the counter values (while the other side can only change control-states).

1 Introduction

Markov decision processes (MDPs) [17, 14] are a formal model for games on directed graphs, where certain decisions are taken by a strategic player (a.k.a. Player 1, or controller) while others are taken randomly (a.k.a. by nature, or the environment) according to pre-defined probability distributions. MDPs are thus a subclass of general 2-player stochastic games, and they are equivalent to 1.5-player games in the terminology of [10]. They are also called “games against nature”.

A run of the MDP consists of a sequence of visited states and transitions on the graph. Properties of the system are expressed via properties of the induced runs. The most basic objectives are reachability (is a certain (set of) control-state(s) eventually visited?) and Büchi objectives (is a certain (set of) control-state(s) visited infinitely often?).

^{*} Technical Report EDI-INF-RR-1422 of the School of Informatics at the University of Edinburgh, UK. (<http://www.inf.ed.ac.uk/publications/report/>). Full version (including proofs) of material presented at FoSSaCS 2016 (Eindhoven, The Netherlands). arXiv.org 1512.08824 - CC BY 4.0.

^{**} Supported by UPMARC, Uppsala Programming for Multicore Architectures Research Center.

^{***} Supported by EPSRC grant EP/M027651/1.

[†] Supported by the MIUR-PRIN project CINA and the EU ARTEMIS Joint Undertaking under grant agreement no. 332933 (HoliDes).

Since a strategy of Player 1 induces a probability distribution of runs of the MDP, the objective of an MDP is defined in terms of this distribution, e.g., if the probability of satisfying a reachability/Büchi objective is at least a given constant. The special case where this constant is 1 is a key example of a qualitative objective. Here one asks whether Player 1 has a strategy that achieves an objective surely (all runs satisfy the property) or almost-surely (the probability of the runs satisfying the property is 1).

Most classical work on algorithms for MDPs and stochastic games has focused on finite-state systems (e.g., [14, 19, 11]), but more recently several classes of infinite-state systems have been considered as well. For instance, MDPs and stochastic games on infinite-state probabilistic recursive systems (i.e., probabilistic pushdown automata with unbounded stacks) [13] and on one-counter systems [7, 6] have been studied. Another infinite-state probabilistic model, which is incomparable to recursive systems, is a suitable probabilistic extension of Vector Addition Systems with States (VASS; a.k.a. Petri nets), which have a finite number of unbounded counters holding natural numbers.

Our contribution. We study the decidability of probability-1 qualitative reachability and Büchi objectives for infinite-state MDPs that are induced by suitable probabilistic extensions of VASS that we call VASS-MDPs. (Most quantitative objectives in probabilistic VASS are either undecidable, or the solution is at least not effectively expressible in $(\mathbb{R}, +, *, \leq)$ [2].) It is easy to show that, for general VASS-MDPs, even the simplest of these problems, (almost) sure reachability, is undecidable. Thus we consider two monotone subclasses: 1-VASS-MDPs and P-VASS-MDPs. In 1-VASS-MDPs, only Player 1 can modify counter values while the probabilistic player can only change control-states, whereas for P-VASS-MDPs it is vice-versa. Still these two models induce infinite-state MDPs. Unlike for finite-state MDPs, it is possible that the value of the MDP, in the game theoretic sense, is 1, even though there is no single strategy that achieves value 1. For example, there can exist a family of strategies σ_ϵ for every $\epsilon > 0$, where playing σ_ϵ ensures a probability $\geq 1 - \epsilon$ of reaching a given target state, but no strategy ensures probability 1. In this case, one says that the reachability property holds limit-surely, but not almost-surely (i.e., unlike in finite-state MDPs, almost-surely and limit-surely do not coincide in infinite-state MDPs).

We show that even for P-VASS-MDPs, all sure/almost-sure/limit-sure reachability/Büchi problems are still undecidable. However, in the deadlock-free subclass of P-VASS-MDPs, the sure reachability/Büchi problems become decidable (while the other problems remain undecidable). In contrast, for 1-VASS-MDPs, the sure/almost-sure/limit-sure reachability problem and the sure/almost-sure Büchi problem are decidable.

Our decidability results rely on two different techniques. For the sure and almost sure problems, we prove that we can reduce them to the model-checking problem over VASS of a restricted fragment of the modal μ -calculus that has been proved to be decidable in [3]. For the limit-sure reachability problem in 1-VASS-MDP, we use an algorithm which at each iteration reduces the dimension of the considered VASS while preserving the limit-sure reachability properties.

Although we do not consider the class of qualitative objectives referring to the probability of (repeated) reachability being strictly greater than 0, we observe that reachability on VASS-MDPs in such a setting is equivalent to reachability on standard VASS (though this correspondence does not hold for repeated reachability).

Outline. In Section 2 we define basic notations and how VASS induce Markov decision processes. In Sections 3 and 4 we consider verification problems for P-VASS-MDP and 1-VASS-MDP, respectively. In Section 5 we summarize the decidability results (Table 1) and outline future work.

2 Models and verification problems

Let \mathbb{N} (resp. \mathbb{Z}) denote the set of nonnegative integers (resp. integers). For two integers i, j such that $i \leq j$ we use $[i..j]$ to represent the set $\{k \in \mathbb{Z} \mid i \leq k \leq j\}$. Given a set X and $n \in \mathbb{N} \setminus \{0\}$, X^n is the set of n -dimensional vectors with values in X . We use $\mathbf{0}$ to denote the vector such that $\mathbf{0}(i) = 0$ for all $i \in [1..n]$. The classical order on \mathbb{Z}^n is denoted \leq and is defined by $\mathbf{v} \leq \mathbf{w}$ if and only if $\mathbf{v}(i) \leq \mathbf{w}(i)$ for all $i \in [1..n]$. We also define the operation $+$ over n -dimensional vectors of integers in the classical way (i.e., for $\mathbf{v}, \mathbf{v}' \in \mathbb{Z}^n$, $\mathbf{v} + \mathbf{v}'$ is defined by $(\mathbf{v} + \mathbf{v}')(i) = \mathbf{v}(i) + \mathbf{v}'(i)$ for all $i \in [1..n]$). Given a set S , we use S^* (respectively S^ω) to denote the set of finite (respectively infinite) sequences of elements of S . We now recall the notion of well-quasi-ordering (which we abbreviate as wqo). A quasi-order (A, \preceq) is a wqo if for every infinite sequence of elements a_1, a_2, \dots in A , there exist two indices $i < j$ such that $a_i \preceq a_j$. For $n > 0$, (\mathbb{N}^n, \leq) is a wqo. Given a set A with an ordering \preceq and a subset $B \subseteq A$, the set B is said to be *upward closed* in A if $a_1 \in B$, $a_2 \in A$ and $a_1 \preceq a_2$ implies $a_2 \in B$.

2.1 Markov decision processes

A probability distribution on a countable set X is a function $f : X \mapsto [0, 1]$ such that $\sum_{x \in X} f(x) = 1$. We use $\mathcal{D}(X)$ to denote the set of all probability distributions on X . We first recall the definition of Markov decision processes.

Definition 1 (MDPs). A Markov decision process (MDP) M is a tuple $\langle C, C_1, C_P, A, \rightarrow, p \rangle$ where: C is a countable set of configurations partitioned into C_1 and C_P (that is $C = C_1 \cup C_P$ and $C_1 \cap C_P = \emptyset$); A is a set of actions; $\rightarrow \subseteq C \times A \times C$ is a transition relation; $p : C_P \mapsto \mathcal{D}(C)$ is a partial function which assigns to some configurations in C_P probability distributions on C such that $p(c)(c') > 0$ if and only if $c \xrightarrow{a} c'$ for some $a \in A$.

Note that our definition is equivalent as seeing MDPs as games played between a nondeterministic player (Player 1) and a probabilistic player (Player P). The set C_1 contains the nondeterministic configurations (or configurations of Player 1) and the set C_P contains the probabilistic configurations (or configurations of Player P). Given two configurations c, c' in C , we write $c \rightarrow c'$ whenever

there exists $a \in A$ such that $c \xrightarrow{a} c'$. We will say that a configuration $c \in C$ is a *deadlock* if there does not exist $c' \in C$ such that $c \rightarrow c'$. We use C_1^{df} (resp. C_P^{df}), to denote the configurations of Player 1 (resp. of Player P) which are not a deadlock (*df* stands here for deadlock free).

A *play* of the MDP $M = \langle C, C_1, C_P, A, \rightarrow, p \rangle$ is either an infinite sequence of the form $c_0 \xrightarrow{a_0} c_1 \xrightarrow{a_1} c_2 \cdots$ or a finite sequence $c_0 \xrightarrow{a_0} c_1 \xrightarrow{a_1} c_2 \cdots \xrightarrow{a_{k-1}} c_k$. We call the first kind of play an *infinite play*, and the second one a *finite play*. A play is said to be maximal whenever it is infinite or it ends in a deadlock configuration. These latter plays are called *deadlocked plays*. We use Ω to denote the set of maximal plays. For a finite play $\rho = c_0 \xrightarrow{a_0} c_1 \xrightarrow{a_1} c_2 \cdots \xrightarrow{a_{k-1}} c_k$, let $c_k = \text{last}(\rho)$. We use Ω_1^{df} to denote the set of finite plays ρ such that $\text{last}(\rho) \in C_1^{df}$.

A *strategy* for Player 1 is a function $\sigma : \Omega_1^{df} \mapsto C$ such that, for all $\rho \in \Omega_1^{df}$ and $c \in C$, if $\sigma(\rho) = c$ then $\text{last}(\rho) \rightarrow c$. Intuitively, given a finite play ρ , which represents the history of the game so far, the strategy represents the choice of Player 1 among the different possible successor configurations from $\text{last}(\rho)$. We use Σ to denote the set of all strategies for Player 1. Given a strategy $\sigma \in \Sigma$, an infinite play $c_0 \xrightarrow{a_0} c_1 \xrightarrow{a_1} c_2 \cdots$ *respects* σ if for every $k \in \mathbb{N}$, we have that if $c_k \in C_1$ then $c_{k+1} = \sigma(c_0 \xrightarrow{a_0} c_1 \xrightarrow{a_1} c_2 \cdots c_k)$ and if $c_k \in C_P$ then $p(c_k)(c_{k+1}) > 0$. We define finite plays that respect σ similarly. Let $\text{Plays}(M, c, \sigma) \subseteq \Omega$ be the set of all maximal plays of M that start from c and that respect σ .

Note that once a starting configuration $c_0 \in C$ and a strategy σ have been chosen, the MDP is reduced to an ordinary stochastic process. We define an event $\mathcal{A} \subseteq \Omega$ as a measurable set of plays and we use $\mathbb{P}(M, c, \sigma, \mathcal{A})$ to denote the probability of event \mathcal{A} starting from $c \in C$ under strategy σ . The notation $\mathbb{P}^+(M, c, \mathcal{A})$ will be used to represent the maximal probability of event \mathcal{A} starting from c which is defined as $\mathbb{P}^+(M, c, \mathcal{A}) = \sup_{\sigma \in \Sigma} \mathbb{P}(M, c, \sigma, \mathcal{A})$.

2.2 VASS-MDPs

Probabilistic Vector Addition Systems with States have been studied, e.g., in [2]. Here we extend this model with non-deterministic choices made by a controller. We call this new model VASS-MDPs. We first recall the definition of Vector Addition Systems with States.

Definition 2 (Vector Addition System with States). *For $n > 0$, an n -dimensional Vector Addition System with States (VASS) is a tuple $S = \langle Q, T \rangle$ where Q is a finite set of control states and $T \subseteq Q \times \mathbb{Z}^n \times Q$ is the transition relation labelled with vectors of integers.*

In the sequel, we will not always make precise the dimension of the considered VASS. Configurations of a VASS are pairs $\langle q, \mathbf{v} \rangle \in Q \times \mathbb{N}^n$. Given a configuration $\langle q, \mathbf{v} \rangle$ and a transition $t = \langle q, \mathbf{z}, q' \rangle$ in T , we will say that t is *enabled* at $\langle q'', \mathbf{v} \rangle$, if $q = q''$ and $\mathbf{v} + \mathbf{z} \geq \mathbf{0}$. Let then $\text{En}(q, \mathbf{v})$ be the set $\{t \in T \mid t \text{ is enabled at } \langle q, \mathbf{v} \rangle\}$. In case the transition $t = \langle q, \mathbf{z}, q' \rangle$ is enabled at $\langle q, \mathbf{v} \rangle$, we define $t(q, \mathbf{v}) = \langle q', \mathbf{v}' \rangle$ where $\mathbf{v}' = \mathbf{v} + \mathbf{z}$. An n -dimensional VASS S induces a labelled transition system $\langle C, T, \rightarrow \rangle$ where $C = Q \times \mathbb{N}^n$ is the set of configurations and the transition

relation $\rightarrow \subseteq C \times T \times C$ is defined as follows: $\langle q, \mathbf{v} \rangle \xrightarrow{t} \langle q', \mathbf{v}' \rangle$ iff $\langle q', \mathbf{v}' \rangle = t(q, \mathbf{v})$. VASS are sometimes seen as programs manipulating integer variables, a.k.a. counters. When a transition of a VASS changes the i -th value of a vector \mathbf{v} , we will sometimes say that it modifies the value of the i -th counter. We show now in which manner we add probability distributions to VASS.

Definition 3 (VASS-MDP). *A VASS-MDP is a tuple $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ where $\langle Q, T \rangle$ is a VASS for which the set of control states Q is partitioned into Q_1 and Q_P , and $\tau : T \mapsto \mathbb{N} \setminus \{0\}$ is a partial function assigning to each transition a weight which is a positive natural number.*

Nondeterministic (resp. probabilistic) choices are made from control states in Q_1 (resp. Q_P). The subset of transitions from control states of Q_1 (resp. control states of Q_P) is denoted by T_1 (resp. T_P). Hence $T = T_1 \cup T_P$ with $T_1 \subseteq Q_1 \times \mathbb{Z}^n \times Q$ and $T_P \subseteq Q_P \times \mathbb{Z}^n \times Q$. A VASS-MDP $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ induces an MDP $M_S = \langle C, C_1, C_P, T, \rightarrow, p \rangle$ where: $\langle C, T, \rightarrow \rangle$ is the labelled transition system associated with the VASS $\langle Q, T \rangle$; $C_1 = Q_1 \times \mathbb{N}^n$ and $C_P = Q_P \times \mathbb{N}^n$; and for all $c \in C_P^{df}$ and $c' \in C$, if $c \rightarrow c'$, the probability of going from c to c' is defined by $p(c)(c') = (\sum_{\{t | t(c) = c'\}} \tau(t)) / (\sum_{t \in \text{En}(c)} \tau(t))$, whereas if $c \not\rightarrow c'$, we have $p(c)(c') = 0$. Note that the MDP M_S is well-defined: when defining $p(c)(c')$ in the case $c \rightarrow c'$, there exists at least one transition in $\text{En}(c)$ and consequently the sum $\sum_{t \in \text{En}(c)} \tau(t)$ is never equal to 0. Also, we could have restricted the weights to be assigned only to transitions leaving from a control state in Q_P since we do not take into account the weights assigned to the other transitions. A VASS-MDP is deadlock free if its underlying VASS is deadlock free.

Finally, as in [18] or [3], we will see that to gain decidability it is useful to restrict the power of the nondeterministic player or of the probabilistic player by restricting their ability to modify the counters' values and hence letting them only choose a control location. This leads to the two following definitions: a *P-VASS-MDP* is a VASS-MDP $\langle Q, Q_1, Q_P, T, \tau \rangle$ such that for all $\langle q, \mathbf{z}, q' \rangle \in T_1$, we have $\mathbf{z} = \mathbf{0}$ and a *1-VASS-MDP* is a VASS-MDP $\langle Q, Q_1, Q_P, T, \tau \rangle$ such that for all $\langle q, \mathbf{z}, q' \rangle \in T_P$, we have $\mathbf{z} = \mathbf{0}$. In other words, in a P-VASS-MDP, Player 1 cannot change the counter values when taking a transition and, in a 1-VASS-MDP, it is Player P which cannot perform such an action.

2.3 Verification problems for VASS-MDPs

We consider qualitative verification problems for VASS-MDPs, taking as objectives control-state reachability and repeated reachability. To simplify the presentation, we consider a single target control-state $q_F \in Q$. However, our positive decidability results easily carry over to sets of target control-states (while the negative ones trivially do). Note however, that asking to reach a *fixed target configuration* like $\langle q_F, \mathbf{0} \rangle$ is a very different problem (cf. [2]).

Let $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ be a VASS-MDP and M_S its associated MDP. Given a control state $q_F \in Q$, we denote by $\llbracket \Diamond q_F \rrbracket$ the set of infinite plays $c_0 \cdot c_1 \cdots$ and deadlocked plays $c_0 \cdots c_l$ of M_S for which there exists an index

$k \in \mathbb{N}$ such that $c_k = \langle q_F, \mathbf{v} \rangle$ for some $\mathbf{v} \in \mathbb{N}^n$. Similarly, $\llbracket \Box \Diamond q_F \rrbracket$ characterizes the set of infinite plays $c_0 \cdot c_1 \cdot \dots$ of M_S for which the set $\{i \in \mathbb{N} \mid c_i = \langle q_F, \mathbf{v} \rangle \text{ for some } \mathbf{v} \in \mathbb{N}^n\}$ is infinite. Since M_S is an MDP with a countable number of configurations, we know that the sets of plays $\llbracket \Diamond q_F \rrbracket$ and $\llbracket \Box \Diamond q_F \rrbracket$ are measurable (for more details see for instance [4]), and are hence events for M_S . Given an initial configuration $c_0 \in Q \times \mathbb{N}^n$ and a control-state $q_F \in Q$, we consider the following questions:

1. The *sure reachability problem*: Does there exist a strategy $\sigma \in \Sigma$ such that $\text{Plays}(M_S, c_0, \sigma) \subseteq \llbracket \Diamond q_F \rrbracket$?
2. The *almost-sure reachability problem*: Does there exist a strategy $\sigma \in \Sigma$ such that $\mathbb{P}(M_S, c_0, \sigma, \llbracket \Diamond q_F \rrbracket) = 1$?
3. The *limit-sure reachability problem*: Does $\mathbb{P}^+(M_S, c_0, \llbracket \Diamond q_F \rrbracket) = 1$?
4. The *sure repeated reachability problem*: Does there exist a strategy $\sigma \in \Sigma$ such that $\text{Plays}(M_S, c_0, \sigma) \subseteq \llbracket \Box \Diamond q_F \rrbracket$?
5. The *almost-sure repeated reachability problem*: Does there exist a strategy $\sigma \in \Sigma$ such that $\mathbb{P}(M_S, c_0, \sigma, \llbracket \Box \Diamond q_F \rrbracket) = 1$?
6. The *limit-sure repeated reachability problem*: Does $\mathbb{P}^+(M_S, c_0, \llbracket \Box \Diamond q_F \rrbracket) = 1$?

Note that sure reachability implies almost-sure reachability, which itself implies limit-sure reachability, but not vice-versa, as shown by the counterexamples in Figure 1 (see also [7]). The same holds for repeated reachability. For the sure problems, probabilities are not taken into account, and thus these problems can be interpreted as the answer to a two player reachability game played on the transition system of S . Such games have been studied for instance in [18, 1, 3]. Finally, VASS-MDPs subsume deadlock-free VASS-MDPs and thus decidability (resp. undecidability) results carry over to the smaller (resp. larger) class.




Fig. 1. Two 1-dimensional VASS-MDPs. The circles (resp. squares) are the control states of Player 1 (resp. Player P). All transitions have the same weight 1. From $\langle q_0, 0 \rangle$, the state q_F is reached almost-surely, but not surely, due to the possible run with an infinite loop at q_0 (which has probability zero). From $\langle q_1, 0 \rangle$, the state q_F can be reached limit-surely (by a family of strategies that repeats the loop at q_1 more and more often), but not almost-surely (or surely), since every strategy has a chance of getting stuck at state q_2 with counter value zero.

2.4 Undecidability in the general case

It was shown in [1] that the sure reachability problem is undecidable for (2-dimensional) two player VASS. From this we can deduce that the sure reachability problem is undecidable for VASS-MDPs. We now present a similar proof to show the undecidability of the almost-sure reachability problem for VASS-MDPs.

For all of our undecidability results we use reductions from the undecidable control-state reachability problem for Minsky machines. A Minsky machine is a tuple $\langle Q, T \rangle$ where Q is a finite set of states and T is a finite set of transitions manipulating two counters, say x_1 and x_2 . Each transition is a triple of the form $\langle q, x_i = 0?, q' \rangle$ (counter x_i is tested for 0) or $\langle q, x_i := x_i + 1, q' \rangle$ (counter x_i is incremented) or $\langle q, x_i := x_i - 1, q' \rangle$ (counter x_i is decremented) where $q, q' \in Q$. Configurations of a Minsky machine are triples in $Q \times \mathbb{N} \times \mathbb{N}$. The transition relation \Rightarrow between configurations of the Minsky machine is then defined in the obvious way. Given an initial state q_I and a final state q_F , the control-state reachability problem asks whether there exists a sequence of configurations $\langle q_I, 0, 0 \rangle \Rightarrow \langle q_1, v_1, v'_1 \rangle \Rightarrow \dots \Rightarrow \langle q_k, v_k, v'_k \rangle$ with $q_k = q_F$. This problem is known to be undecidable [16]. W.l.o.g. we assume that Minsky machines are deadlock-free and deterministic (i.e., each configuration has always a unique successor) and that the only transition leaving q_F is of the form $\langle q_F, x_1 := x_1 + 1, q_F \rangle$.




Fig. 2. Encoding $\langle q_1, x_1 := x_1 + 1, q_2 \rangle$ and $\langle q_3, x_2 := x_2 - 1, q_4 \rangle$ and $\langle q_5, x_1 = 0?, q_6 \rangle$

We now show how to reduce the control-state reachability problem to the almost-sure and limit-sure reachability problems in deadlock-free VASS-MDPs. From a Minsky machine, we construct a deadlock-free 2-dimensional VASS-MDP for which the control states of Player 1 are exactly the control states of the Minsky machine. The encoding is presented in Figure 2 where the circles (resp. squares) are the control states of Player 1 (resp. Player P), and for each edge the corresponding weight is 1. The state \perp is an absorbing state from which the unique outgoing transition is a self loop that does not affect the values of the counters. This encoding allows us to deduce our first result.

Theorem 1. *The sure, almost-sure and limit-sure (repeated) reachability problems are undecidable problems for 2-dimensional deadlock-free VASS-MDPs.*

In the special case of 1-dimensional VASS-MDPs, the sure and almost-sure reachability problems are decidable [7].

2.5 Model-checking μ -calculus on single-sided VASS

It is well-known that there is a strong connection between model-checking branching time logics and games, and in our case we have in fact undecidability results for simple reachability games played on a VASS and for the model-checking of VASS with expressive branching-time logics [12]. However for this latter point,

decidability can be regained by imposing some restrictions on the VASS structure [3] as we will now recall. We say that a VASS $\langle Q, T \rangle$ is (Q_1, Q_2) -single-sided iff Q_1 and Q_2 represents a partition of the set of states Q such that for all transitions $\langle q, \mathbf{z}, q' \rangle$ in T with $q \in Q_2$, we have $\mathbf{z} = \mathbf{0}$; in other words only the transitions leaving a state from Q_1 are allowed to change the values of the counters. In [3], it has been shown that, thanks to a reduction to games played on a single-sided VASS with parity objectives, a large fragment of the μ -calculus called L_μ^{sv} has a decidable model-checking problem over single-sided VASS. The idea of this fragment is that the “always” operator \Box is guarded with a predicate enforcing the current control states to belong to Q_2 . Formally, the syntax of L_μ^{sv} for (Q_1, Q_2) -single-sided VASS is given by the following grammar: $\phi ::= q \mid X \mid \phi \wedge \phi \mid \phi \vee \phi \mid \Diamond \phi \mid Q_2 \wedge \Box \phi \mid \mu X. \phi \mid \nu X. \phi$, where Q_2 stands for the formula $\bigvee_{q \in Q_2} q$ and X belongs to a set of variables \mathcal{X} . The semantics of L_μ^{sv} is defined as usual: it associates to a formula ϕ and to an environment $\varepsilon : \mathcal{X} \rightarrow 2^C$ a subset of configurations $\llbracket \phi \rrbracket_\varepsilon$. We use ε_0 to denote the environment which assigns the empty set to any variable. Given an environment ε , a variable $X \in \mathcal{X}$ and a subset of configurations C , we use $\varepsilon[X := C]$ to represent the environment ε' which is equal to ε except on the variable X , where we have $\varepsilon'(X) = C$. Finally the notation $\llbracket \phi \rrbracket$ corresponds to the interpretation $\llbracket \phi \rrbracket_{\varepsilon_0}$.

The problem of model-checking single-sided VASS with L_μ^{sv} can then be defined as follows: given a single-sided VASS $\langle Q, T \rangle$, an initial configuration c_0 and a formula ϕ of L_μ^{sv} , do we have $c_0 \in \llbracket \phi \rrbracket$?

Theorem 2. [3] *Model-checking single-sided VASS wrt. L_μ^{sv} is decidable.*

3 Verification of P-VASS-MDPs

In [3] it is proved that parity games played on a single-sided deadlock-free VASS are decidable (this entails the decidability of model checking L_μ^{sv} over single-sided VASS). We will see here that in the case of P-VASS-MDPs, in which only the probabilistic player can modify the counters, the decidability status depends on the presence of deadlocks in the system.

3.1 Undecidability in presence of deadlocks

We point out that the reduction presented in Figure 2 to prove Theorem 1 does not carry over to P-VASS-MDPs, because in that construction both players have the ability to change the counter values. However, it is possible to perform a similar reduction leading to the undecidability of verification problems for P-VASS-MDPs, the main difference being that we crucially exploit the fact that the P-VASS-MDP can contain deadlocks.

We now explain the idea behind our encoding of Minsky machines into P-VASS-MDPs. Intuitively, Player 1 chooses a transition of the Minsky machine to simulate, anticipating the modification of the counters values, and Player P is then in charge of performing the change. If Player 1 chooses a transition with a

decrement and the accessed counter value is actually 0, then Player P will be in a deadlock state and consequently the desired control state will not be reached. Furthermore, if Player 1 decides to perform a zero-test when the counter value is strictly positive, then Player P is able to punish this choice by entering a deadlock state. Similarly to the proof of Theorem 1, Player P can test if the value of the counter is strictly greater than 0 by decrementing it. The encoding of the Minsky machine is presented in Figure 3. Note that no outgoing edge of Player 1's states changes the counter values. Furthermore, we see that Player P reaches the control state \perp if and only if Player 1 chooses to take a transition with a zero-test when the value of the tested counter is not equal to 0. Note that, with the encoding of the transition $\langle q_3, x_2 := x_2 - 1, q_4 \rangle$, when Player P is in the control state between q_3 and q_4 , it can be in a deadlock if the value of the second counter is not positive. In the sequel we will see that in P-VASS-MDP without deadlocks the sure reachability problem becomes decidable.




Fig. 3. Encoding $\langle q_1, x_1 := x_1 + 1, q_2 \rangle$ and $\langle q_3, x_2 := x_2 - 1, q_4 \rangle$ and $\langle q_5, x_1 = 0?, q_6 \rangle$

From this encoding we deduce the following result.

Theorem 3. *The sure, almost sure and limit sure (repeated) reachability problems are undecidable for 2-dimensional P-VASS-MDPs.*

3.2 Sure (repeated) reachability in deadlock-free P-VASS-MDPs

Unlike in the case of general P-VASS-MDPs, we will see that the sure (repeated) reachability problem is decidable for deadlock-free P-VASS-MDPs. Let $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ be a deadlock-free P-VASS-MDP, $M_S = (C, C_1, C_P, \rightarrow, p)$ its associated MDP and $q_F \in Q$ a control state. Note that because the P-VASS-MDP S is deadlock free, Player P cannot take the play to a deadlock to avoid the control state q_F , but he has to deal only with infinite plays. Since S is a P-VASS-MDP, the VASS $\langle Q, T \rangle$ is (Q_P, Q_1) -single-sided. In [18, 1], it has been shown that control-state reachability games on deadlock-free single-sided VASS are decidable, and this result has been extended to parity games in [3]. This implies the decidability of sure (repeated) reachability in deadlock-free P-VASS-MDPs. However, to obtain a generic way of verifying these systems, we construct

a formula of L_μ^{sv} that characterizes the sets of winning configurations and use then the result of Theorem 2. Let V_S^P be the set of configurations from which the answer to the sure reachability problem (with q_F as state to be reached) is negative, i.e., $V_S^P = \{c \in C \mid \nexists \sigma \in \Sigma \text{ s.t. } \text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Diamond q_F \rrbracket\}$ and similarly let $W_S^P = \{c \in C \mid \nexists \sigma \in \Sigma \text{ s.t. } \text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Box \Diamond q_F \rrbracket\}$. The next lemma relates these two sets with a formula of L_μ^{sv} (where Q_P corresponds to the formula $\bigvee_{q \in Q_P}$ and Q_1 corresponds to the formula $\bigvee_{q \in Q_1}$).

Lemma 1.

- $V_S^P = \llbracket \nu X. (\bigvee_{q \in Q \setminus \{q_F\}} q) \wedge (Q_1 \vee \Diamond X) \wedge (Q_P \vee (Q_1 \wedge \Box X)) \rrbracket$.
- $W_S^P = \llbracket \mu Y. \nu X. ((\bigvee_{q \in Q \setminus \{q_F\}} q) \wedge (Q_1 \vee \Diamond X) \wedge (Q_P \vee (Q_1 \wedge \Box X))) \vee (q_F \wedge Q_P \wedge \Diamond Y) \vee (q_F \wedge Q_1 \wedge \Box Y) \rrbracket$

Note that we use $(Q_P \vee (Q_1 \wedge \Box X))$ instead of $(Q_P \vee \Box X)$ so that the formulae are in the guarded fragment of the μ -calculus. Since the two formulae belong to L_μ^{sv} for the (Q_P, Q_1) -single-sided VASS S , decidability follows directly from Theorem 2.

Theorem 4. *The sure reachability and repeated reachability problem are decidable for deadlock free P-VASS-MDPs.*

3.3 Almost-sure and limit-sure reachability in deadlock-free P-VASS-MDPs

We have seen that, unlike for the general case, the sure reachability and sure repeated reachability problems are decidable for deadlock free P-VASS-MDPs, with deadlock freeness being necessary to obtain decidability. For the corresponding almost-sure and limit-sure problems we now show undecidability, again using a reduction from the reachability problem for two counter Minsky machines, as shown in Figure 4. The main difference with the construction used for the proof of Theorem 3 lies in the addition of a self-loop in the encoding of the transitions for decrementing a counter, in order to avoid deadlocks. If Player 1, from a configuration $\langle q_3, \mathbf{v} \rangle$, chooses the transition $\langle q_3, x_2 := x_2 - 1, q_4 \rangle$ which decrements the second counter, then the probabilistic state with the self-loop is entered, and there are two possible cases: if $\mathbf{v}(2) > 0$ then the probability of staying forever in this loop is 0 and the probability of eventually going to state q_4 is 1; on the other hand, if $\mathbf{v}(2) = 0$ then the probability of staying forever in the self-loop is 1, since the other transition that leaves the state of Player P and which performs the decrement on the second counter effectively is not available. Note that such a construction does not hold in the case of sure reachability, because the path that stays forever in the loop is a valid path.

This allows us to deduce the following result for deadlock free P-VASS-MDPs.

Theorem 5. *The almost-sure and limit-sure (repeated) reachability problems are undecidable for 2-dimensional deadlock-free P-VASS-MDPs.*




Fig. 4. Encoding $\langle q_1, x_1 := x_1 + 1, q_2 \rangle$ and $\langle q_3, x_2 := x_2 - 1, q_4 \rangle$ and $\langle q_5, x_1 = 0?, q_6 \rangle$

4 Verification of 1-VASS-MDPs

In this section, we will provide decidability results for the subclass of 1-VASS-MDPs. As for deadlock-free P-VASS-MDPs, the proofs for sure and almost-sure problems use the decidability of L_μ^{sv} over single-sided VASS, whereas the technique used to show decidability of limit-sure reachability is different.

4.1 Sure problems in 1-VASS-MDPs

First we show that, unlike for P-VASS-MDPs, deadlocks do not matter for 1-VASS-MDPs. The idea is that in this case, if the deadlock is in a probabilistic configuration, it means that there is no outgoing edge (because of the property of 1-VASS-MDPs), and hence one can add an edge to a new absorbing state, and the same can be done for the states of Player 1. Such a construction does not work for P-VASS-MDPs, because in that case deadlocks in probabilistic configurations may depend on the counter values, and not just on the current control-state.

Lemma 2. *The sure (resp. almost sure, resp. limit sure) (repeated) reachability problem for 1-VASS-MDPs reduces to the sure (resp. almost sure, resp. limit-sure) (repeated) reachability problem for deadlock-free 1-VASS-MDPs.*

Hence in the sequel we will consider only deadlock-free 1-VASS-MDPs. Let $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ be a deadlock-free 1-VASS-MDP. For what concerns the sure (repeated) reachability problems we can directly reuse the results from Lemma 1 and then show that the complement formulae of the ones expressed in this lemma belong to L_μ^{sv} for the (Q_1, Q_P) -single-sided VASS $\langle Q, T \rangle$ (in fact the correctness of these two lemmas did not depend on the fact that we were considering P-VASS-MDPs). Theorem 2 allows us to retrieve the decidability results already expressed in [18] (for sure reachability) and [3] (for sure repeated reachability).

Theorem 6. *The sure (repeated) reachability problem is decidable for 1-VASS-MDPs.*

4.2 Almost-sure problems in 1-VASS-MDPs

We now move to the case of almost-sure problems in 1-VASS-MDPs. We consider a deadlock free 1-VASS-MDP $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ and its associated MDP $M_S = \langle C, C_1, C_P, \rightarrow, p \rangle$. We will see that, unlike for P-VASS-MDPs, it is here also possible to characterize by formulae of L_μ^{sv} the two following sets: $V_{AS}^1 = \{c \in C \mid \exists \sigma \in \Sigma \text{ such that } \mathbb{P}(M_S, c, \sigma, \llbracket \Diamond q_F \rrbracket) = 1\}$ and $W_{AS}^1 = \{c \in C \mid \exists \sigma \in \Sigma \text{ such that } \mathbb{P}(M_S, c, \sigma, \llbracket \Box \Diamond q_F \rrbracket) = 1\}$, i.e. the set of configurations from which Player 1 has a strategy to reach the control state q_F , respectively to visit infinitely often q_F , with probability 1.

We begin with introducing the following formula of L_μ^{sv} based on the variables X and Y : $\text{InvPre}(X, Y) = (Q_1 \wedge \Diamond(X \wedge Y)) \vee (\Diamond Y \wedge Q_P \wedge \Box X)$. Note that $\text{InvPre}(X, Y)$ is a formula of L_μ^{sv} for the (Q_1, Q_P) -single-sided VASS $\langle Q, T \rangle$. Intuitively, this formula represents the set of configurations from which (i) Player 1 can make a transition to the set represented by the intersection of the sets characterized by the variables X and Y and (ii) Player P can make a transition to the set Y and cannot avoid making a transition to the set X .

Almost sure reachability. We will now prove that V_{AS}^1 can be characterized by the following formula of L_μ^{sv} : $\nu X. \mu Y. (q_F \vee \text{InvPre}(X, Y))$. Note that a similar result exists for finite-state MDPs, see e.g. [9]; this result in general does not extend to infinite-state MDPs, but in the case of VASS-MDPs it can be applied. Before proving this we need some intermediate results.

We denote by E the set $\llbracket \nu X. \mu Y. (q_F \vee \text{InvPre}(X, Y)) \rrbracket_{\varepsilon_0}$. Since $\nu X. \mu Y. (q_F \vee \text{InvPre}(X, Y))$ is a formula of L_μ^{sv} interpreted over the single-sided VASS $\langle Q, T \rangle$, we can show that E is an upward-closed set. We now need another lemma which states that there exists $N \in \mathbb{N}$ and a strategy for Player 1 such that, from any configuration of E , Player 1 can reach the control state q_F in less than N steps and Player P cannot take the play outside of E . The fact that we can bound the number of steps is crucial to show that $\llbracket \nu X. \mu Y. (q_F \vee \text{InvPre}(X, Y)) \rrbracket_{\varepsilon_0}$ is equal to V_{AS}^1 . For infinite-state MDPs where this property does not hold, our techniques do not apply.

Lemma 3. *There exists $N \in \mathbb{N}$ and a strategy σ of Player 1 such that for all $c \in E$, there exists a play $c \cdot c_1 \cdot c_2 \cdot \dots$ in $\text{Plays}(M_S, c, \sigma)$ satisfying the three following properties: (1) there exists $0 \leq i \leq N$ such that $c_i \in \llbracket q_F \rrbracket$; (2) for all $0 \leq j \leq i$, $c_j \in E$; (3) for all $0 \leq j \leq i$, if $c_j \in C_P$ then for all $c'' \in C$ such that $c_j \rightarrow c''$, we have $c'' \in E$.*

This previous lemma allows us to characterize V_{AS}^1 with a formula of L_μ^{sv} . The proof of the following result uses the fact that the number of steps is bounded, and also the fact that the sets described by closed L_μ^{sv} formulae are upward-closed. This makes the fixpoint iteration terminate in a finite number of steps.

Lemma 4. $V_{AS}^1 = \llbracket \nu X. \mu Y. (q_F \vee \text{InvPre}(X, Y)) \rrbracket$.

Since $\langle Q, T \rangle$ is (Q_1, Q_P) -single-sided and since the formula associated to V_{AS}^1 belongs to L_μ^{sv} , from Theorem 2 we deduce the following theorem.

Theorem 7. *The almost-sure reachability problem is decidable for 1-VASS-MDPs.*

Almost sure repeated reachability. For the case of almost sure repeated reachability we reuse the previously introduced formula $\text{InvPre}(X, Y)$. We can perform a reasoning similar to the previous ones and provide a characterization of the set W_{AS}^1 .

Lemma 5. $W_{AS}^1 = \llbracket \nu X. \text{InvPre}(X, \mu Y. (q_F \vee \text{InvPre}(X, Y))) \rrbracket$.

As previously, this allows us to deduce the decidability of the almost sure repeated reachability problem for 1-VASS-MDP.

Theorem 8. *The almost sure repeated reachability problem is decidable for 1-VASS-MDPs.*

4.3 Limit-sure reachability in 1-VASS-MDP

We consider a slightly more general version of the limit-sure reachability problem with a set $X \subseteq Q$ of target states instead of a single state q_F , i.e., the standard case corresponds to $X = \{q_F\}$.

We extend the set of natural numbers \mathbb{N} to $\mathbb{N}_* = \mathbb{N} \cup \{*\}$ by adding an element $*$ $\notin \mathbb{N}$ with $*+j = *-j = *$ and $j < *$ for all $j \in \mathbb{N}$. We consider then the set of vectors \mathbb{N}_*^d . The projection of a vector \mathbf{v} in \mathbb{N}^d by eliminating components that are indexed by a natural number k is defined by $\text{proj}_k(\mathbf{v})(i) = \mathbf{v}(i)$ if $i \neq k$ and $\text{proj}_k(\mathbf{v})(i) = *$ otherwise

Let Q_c represent control-states which are indexed by a color. The coloring functions $\text{col}_i : Q \rightarrow Q_c$ create colored copies of control-states by $\text{col}_i(q) = q_i$.

Given a 1-VASS-MDP $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ of dimension d , an index $k \leq d$ and a color i , the colored projection is defined as:

$$\text{Proj}_k(S, d, i) = \langle \text{col}_i(Q), \text{col}_i(Q_1), \text{col}_i(Q_P), \text{proj}_{k,i}(T), \tau_{k,i} \rangle$$

where $\text{proj}_{k,i}(T) = \{\text{proj}_{k,i}(t) \mid t \in T\}$ is the projection of the set of transitions T and $\text{proj}_{k,i}(t) = \langle \text{col}_i(x), \text{proj}_k(\mathbf{z}), \text{col}_i(y) \rangle$ is the projection of transition $t = \langle x, \mathbf{z}, y \rangle$ obtained by removing component k and coloring the states x and y with color i . The transition weights carry over, i.e., $\tau_{k,i}(t') = \sum \{\tau(t) \mid \text{proj}_{k,i}(t) = t'\}$.

We define the functions $\text{state} : Q \times \mathbb{N}_*^d \rightarrow Q$ and $\text{count} : Q \times \mathbb{N}_*^d \rightarrow \mathbb{N}_*^d$ s.t for a configuration $c_i = \langle q, \mathbf{v} \rangle$, where $q \in Q$ and $\mathbf{v} \in \mathbb{N}_*^d$ we have that $\text{state}(q, \mathbf{v}) = q$ and $\text{count}(q, \mathbf{v}) = \mathbf{v}$. For any two configurations c_1 and c_2 , we write $c_1 \prec c_2$ to denote that $\text{state}(c_1) = \text{state}(c_2)$, and there exists a nonempty set of indexes I where for every $i \in I$, $\text{count}(c_1)(i) < \text{count}(c_2)(i)$, whereas for every index $j \notin I$, $0 < j \leq d$, $\text{count}(c_1)(j) = \text{count}(c_2)(j)$.

Algorithm 1 reduces the dimension of the limit-sure reachability problem for 1-VASS-MDP by a construction resembling the Karp-Miller tree [15]. It takes as

input a 1-VASS-MDP S of some dimension $d > 0$ with a set of target states X . It outputs a new 1-VASS-MDP S' of dimension $d - 1$ and a new set of target states X' such that M_S can limit-surely reach X iff $M_{S'}$ can limit-surely reach X' . In particular, in the base case where $d - 1 = 0$, the new system S' has dimension zero and thus induces a finite-state MDP $M_{S'}$, for which limit-sure reachability of X' coincides with almost-sure reachability of X' , which is known to be decidable in polynomial time. Algorithm 1 starts by exploring all branches of the computation tree of S (and adding them to S' as the so-called initial *uncolored part*) until it encounters a configuration that is either (1) equal to, or (2) strictly larger than a configuration encountered previously on the same branch. In case (1) it just adds a back loop to the point where the configuration was encountered previously. In case (2), it adds a modified copy of S (identified by a unique color) to S' . This so-called colored subsystem is similar to S except that those counters that have strictly increased along the branch are removed. The intuition is that these counters could be pumped to arbitrarily high values and thus present no obstacle to reaching the target. Since the initial uncolored part is necessarily finite (by Dickson's Lemma) and each of the finitely many colored subsystems only has dimension $d - 1$ (since a counter is removed; possibly a different one in different colored subsystems), the resulting 1-VASS-MDP S' has dimension $d - 1$. The set of target states X' is defined as the union of all appearances of states in X in the uncolored part, plus all colored copies of states from X in the colored subsystems.

By Dickson's Lemma, the conditions on line 7 or line 19 of the algorithm must eventually hold on every branch of the explored computation tree. Thus, it will terminate.

Lemma 6. *Algorithm 1 terminates.*

The next lemma states the correctness of Algorithm 1. Let $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ be 1-VASS-MDP of dimension $d > 0$ with initial configuration $c_0 = \langle q_0, \mathbf{v} \rangle$ and $X \subseteq Q$ a set of target states. Let $S' = \langle Q', Q'_1, Q'_P, T', \tau' \rangle$ with initial configuration $c'_0 = \langle q'_0, \mathbf{0} \rangle$ and set of target states $X' \subseteq Q'$ be the $(d - 1)$ dimensional 1-VASS-MDP produced by Algorithm 1. As described above we have the following relation between these two systems.

Lemma 7. $\mathbb{P}^+(M_S, c_0, \llbracket \Diamond X \rrbracket) = 1$ iff $\mathbb{P}^+(M_{S'}, c'_0, \llbracket \Diamond X' \rrbracket) = 1$.

By applying the result of the previous lemma iteratively until we obtain a finite-state MDP, we can deduce the following theorem.

Theorem 9. *The limit-sure reachability problem for 1-VASS-MDP is decidable.*

5 Conclusion and Future Work

Table 1 summarizes our results on the decidability of verification problems for subclasses of VASS-MDP. The exact complexity of most problems is still open.

Algorithm 1 Reducing the dimension of the limit-sure reachability problem.

Require: $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ 1-VASS-MDP, dimension $d > 0$, $c_0 = \langle q_0, \mathbf{v} \rangle \in Q \times \mathbb{N}^d$
 $X \subseteq Q$ - set of target states

Ensure: $S' = \langle Q', Q'_1, Q'_P, T', \tau' \rangle$; $c'_0 = \langle q'_0, \mathbf{0} \rangle$; $X' \subseteq Q'$; $\lambda : Q' \rightarrow ((Q \cup Q_c) \times \mathbb{N}_*^d)$

- 1: $Q' \leftarrow \emptyset$; $Q'_1 \leftarrow \emptyset$; $Q'_P \leftarrow \emptyset$; $T' \leftarrow \emptyset$; $\tau' \leftarrow \emptyset$;
- 2: $\text{new}(q')$; $q'_0 \leftarrow q'$; $\lambda(q') \leftarrow c_0$; $Q' \leftarrow \{q'\}$; $i \leftarrow 0$
- 3: **if** $\text{state}(\lambda(q')) \in Q_1$ **then** $Q'_1 \leftarrow \{q'\}$ **else** $Q'_P \leftarrow \{q'\}$
- 4: $\text{ToExplore} \leftarrow \{q'\}$
- 5: **while** $\text{ToExplore} \neq \emptyset$ **do**
- 6: Pick and remove a $q \in \text{ToExplore}$
- 7: **if** $\exists q'. q'$ is previously on the same branch as q and $\lambda(q') \prec \lambda(q)$ **then**
- 8: get indexes I in which the counter is increasing
- 9: pick and remove the first index k from I
- 10: $i \leftarrow i + 1$; // increase color index
- 11: $\text{new}(q'')$;
- 12: $\lambda(q'') \leftarrow \langle \text{col}_i(\text{state}(\lambda(q))), \text{proj}_k(\text{count}(\lambda(q))) \rangle$
- 13: **if** $\text{state}(\lambda(q)) \in Q_1$ **then** $Q'_1 \leftarrow Q'_1 \cup \{q''\}$ **else** $Q'_P \leftarrow Q'_P \cup \{q''\}$
- 14: $T' \leftarrow T' \cup \{\langle q, \mathbf{0}, q'' \rangle\}$; $\tau'(\langle q, \mathbf{0}, q'' \rangle) = 1$;
- 15: $Q'_1 \leftarrow Q'_1 \cup \text{col}_i(Q_1)$; $Q'_P \leftarrow Q'_P \cup \text{col}_i(Q_P)$; $T' \leftarrow T' \cup \text{proj}_{k,i}(T)$;
- 16: $X' \leftarrow X' \cup \text{col}_i(X)$; $\tau' \leftarrow \tau' \cup \tau_{k,i}$
- 17: **else**
- 18: **for** every $t = \langle x, \mathbf{z}, y \rangle$ in T such that $t \in \text{En}(\lambda(q))$ **do**
- 19: **if** $\exists q'. q'$ is previously on the same branch as q and $t(\lambda(q)) = \lambda(q')$ **then**
- 20: $T' \leftarrow T' \cup \{\langle q, \mathbf{z}, q' \rangle\}$;
- 21: **else**
- 22: $\text{new}(q')$; $\lambda(q') \leftarrow t(\lambda(q))$
- 23: $T' \leftarrow T' \cup \{\langle q, \mathbf{z}, q' \rangle\}$; $\tau'(\langle q, \mathbf{z}, q' \rangle) \leftarrow \tau(t)$
- 24: **if** $\text{state}(\lambda(q')) \in Q_1$ **then** $Q'_1 \leftarrow Q'_1 \cup \{q'\}$ **else** $Q'_P \leftarrow Q'_P \cup \{q'\}$
- 25: **if** $\text{state}(\lambda(q')) \in X$ **then** $X' \leftarrow X' \cup \{q'\}$
- 26: $\text{ToExplore} \leftarrow \text{ToExplore} \cup \{q'\}$
- 27: **end if**
- 28: **end for**
- 29: **end if**
- 30: **end while**

Algorithm 1 relies on Dickson’s Lemma for termination, and the algorithm deciding the model-checking problem of Theorem 2 additionally uses the Valk-Jantzen construction repeatedly. However, all these problems are at least as hard as control-state reachability in VASS, and thus EXPSPACE-hard [12].

The decidability of the limit-sure repeated reachability problem for 1-VASS-MDP is open. A hint of its difficulty is given by the fact that there are instances where the property holds even though a small chance of reaching a deadlock cannot be avoided from any reachable configuration. In particular, a solution would require an analysis of the long-run behavior of multi-dimensional random walks induced by probabilistic VASS. However, these may exhibit strange nonregular behaviors for dimensions ≥ 3 , as described in [8] (Section 5).

	P-VASS-MDP	df P-VASS-MDP	1-VASS-MDP
sure reachability	\times (Thm. 3)	\checkmark (Thm. 4)	\checkmark (Thm. 6)
almost-sure reachability	\times (Thm. 3)	\times (Thm. 5)	\checkmark (Thm. 7)
limit-sure reachability	\times (Thm. 3)	\times (Thm. 5)	\checkmark (Thm. 9)
sure repeated reachability	\times (Thm. 3)	\checkmark (Thm. 4)	\checkmark (Thm. 6)
almost-sure repeated reachability	\times (Thm. 3)	\times (Thm. 5)	\checkmark (Thm. 8)
limit-sure repeated reachability	\times (Thm. 3)	\times (Thm. 5)	Open

Table 1. Decidability of verification problems for P-VASS-MDP, deadlock-free P-VASS-MDP and 1-VASS-MDP. A \checkmark stands for decidable and a \times for undecidable.

References

1. P. A. Abdulla, A. Bouajjani, and J. d'Orso. Monotonic and downward closed games. *Journal of Logic and Computation*, 18(1):153–169, 2008.
2. P. A. Abdulla, N. B. Henda, and R. Mayr. Decisive Markov chains. *Logical Methods in Computer Science*, 3(4), 2007.
3. P. A. Abdulla, R. Mayr, A. Sangnier, and J. Sproston. Solving parity games on integer vectors. In *CONCUR'13*, volume 8052 of *LNCS*, pages 106–120. Springer, 2013.
4. C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
5. N. Bertrand. *Modèles stochastiques pour les pertes de messages dans les protocoles asynchrones et techniques de vérification automatique*. Phd thesis, Laboratoire Spécification et Vérification, ENS Cachan, France, 2006.
6. T. Brázdil, V. Brozek, and K. Etessami. One-counter stochastic games. In *FSTTCS'10*, volume 8 of *LIPICs*, pages 108–119. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2010.
7. T. Brázdil, V. Brozek, K. Etessami, A. Kučera, and D. Wojtczak. One-counter Markov decision processes. In *SODA'10*, pages 863–874. SIAM, 2010.
8. T. Brázdil, S. Kiefer, A. Kučera, and P. Novotný. Long-run average behaviour of probabilistic vector addition systems. In *LICS'15*, pages 44–55. IEEE, 2015.
9. K. Chatterjee, L. de Alfaro, M. Faella, and A. Legay. Qualitative logics and equivalences for probabilistic systems. *Logical Methods in Computer Science*, 5(2), 2009.
10. K. Chatterjee, M. Jurdziński, and T. Henzinger. Simple stochastic parity games. In *CSL'03*, volume 2803 of *LNCS*, pages 100–113. Springer, 2003.
11. A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992.
12. J. Esparza and M. Nielsen. Decidability issues for Petri nets - a survey. *Bulletin of the EATCS*, 52:244–262, 1994.
13. K. Etessami and M. Yannakakis. Recursive Markov decision processes and recursive stochastic games. In *ICALP'05*, volume 3580 of *LNCS*, pages 891–903. Springer, 2005.
14. J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, 1997.
15. R. Karp and R. Miller. Parallel program schemata. *J. Comput. Syst. Sci.*, 3(2):147–195, May 1969.
16. M. L. Minsky. *Computation: finite and infinite machines*. Prentice-Hall, 1967.
17. M. L. Puterman. *Markov Decision Processes*. Wiley, 1994.
18. J.-F. Raskin, M. Samuelides, and L. V. Begin. Games for counting abstractions. In *AVoCS'04*, volume 128(6) of *Electronic Notes in Theoretical Computer Science*, pages 69–85, 2005.
19. L. S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100, 1953.

A Proofs of Section 3

A.1 Proof of Lemma 1

Even if the result of this lemma is quite standard, in order to be consistent we provide the proof, in particular to be sure that the fact that we are dealing with infinite state systems does not harm the reasoning.

Lemma 8. $V_S^P = \llbracket \nu X. (\bigvee_{q \in Q \setminus \{q_F\}} q) \wedge (Q_1 \vee \Diamond X) \wedge (Q_P \vee (Q_1 \wedge \Box X)) \rrbracket$.

Proof. We denote by U the set $\llbracket \nu X. (\bigvee_{q \in Q \setminus \{q_F\}} q) \wedge (Q_1 \vee \Diamond X) \wedge (Q_P \vee (Q_1 \wedge \Box X)) \rrbracket$. We will consider the function $g : 2^C \mapsto 2^C$ such that for each set of configurations $C' \subseteq C$, we have $g(C') = \llbracket \bigvee_{q \in Q \setminus \{q_F\}} q \wedge (Q_1 \vee \Diamond X) \wedge (Q_P \vee (Q_1 \wedge \Box X)) \rrbracket_{\varepsilon_0[X:=C']}$, i.e. the set $\llbracket \bigvee_{q \in Q \setminus \{q_F\}} q \wedge (Q_P \wedge \Diamond X) \wedge (Q_1 \wedge (Q_1 \wedge \Box X)) \rrbracket$ where X is interpreted as C' . Note that U is then the greatest fixpoint of g and hence $U = g(U)$.

We first prove that $U \subseteq V_S^P$. Let c be a configuration in C such that $c \notin V_S^P$. Then there exists a strategy $\sigma \in \Sigma$ such that $\text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Diamond q_F \rrbracket$. We consider such a strategy σ and we reason by contradiction assuming that $c \in U$.

Let us show that since $c \in U$, there exists an infinite play $c_0 \rightarrow c_1 \rightarrow c_2 \rightarrow \dots$ in $\text{Plays}(M_S, c, \sigma)$ such that $c_0 = c$ and $c_i \in U$ for all $i \in \mathbb{N}$. We prove in fact by induction that if $c_0 \rightarrow c_1 \dots \rightarrow c_k$ is a finite play in M_S respecting σ with $c_0 = c$ and such that $c_i \in U$ for $i \in [0..k]$, then there exists $c_{k+1} \in U$ such that $c_0 \rightarrow c_1 \dots \rightarrow c_{k+1}$ is a play in M_S respecting σ . The base case is obvious since $c \in U$. We assume now that $c_0 \rightarrow c_1 \dots \rightarrow c_k$ is a finite play in M_S respecting σ such that $c_i \in U$ for $i \in [0..k]$. Because there is no deadlock, if $c_k \in C_1$ then there exists $c_{k+1} \in C$ satisfying $c_{k+1} = \sigma(c_0 \rightarrow c_1 \dots \rightarrow c_k)$. Furthermore since $c_k \in U$ and $U = g(U)$, we deduce that, for all $c \in C$ such that $c_k \rightarrow c$, we have $c \in U$, and consequently $c_{k+1} \in U$. On the other hand, if $c_k \in C_P$, then since $c_k \in U$, there exists $c_{k+1} \in U$ such that $c_k \rightarrow c_{k+1}$. In both cases, we have that $c_0 \rightarrow c_1 \dots \rightarrow c_k \rightarrow c_{k+1}$ is a play in M_S which respects σ .

We deduce the existence of an infinite play $c_0 \rightarrow c_1 \rightarrow c_2 \rightarrow \dots$ in $\text{Plays}(M_S, c, \sigma)$ such that $c_0 = c$ and $c_i \in U$ for all $i \in \mathbb{N}$. Note that, because $U = g(U)$, we have that $U \subseteq Q \setminus \{q_F\}$. However, because $c_0 \rightarrow c_1 \rightarrow c_2 \rightarrow \dots$ in $\text{Plays}(M_S, c, \sigma)$, we also have that $c_0 \rightarrow c_1 \rightarrow c_2 \rightarrow \dots \in \llbracket \Diamond q_F \rrbracket$, which is a contradiction. Consequently, we have $c \notin U$ and this allows us to conclude that $U \subseteq V_S^P$.

We now prove that $V_S^P \subseteq U$. By the Knaster-Tarski Theorem, since U is the greatest fixpoint of g , we know that $U = \bigcup \{C' \subseteq C \mid C' \subseteq g(C')\}$. It hence suffices to show that $V_S^P \subseteq g(V_S^P)$. Let $c = \langle q, \mathbf{v} \rangle$ be a configuration in V_S^P . First note that by definition of V_S^P , we have $q \neq q_F$. We reason then by a case analysis to prove that $c \in g(V_S^P)$. First assume $c \in C_1$. Then by definition of V_S^P , for all $c' = \langle q', \mathbf{v}' \rangle$ in C satisfying $c \rightarrow c'$, we have $c' \in V_S^P$, otherwise Player 1 would have a strategy to reach q_F from c which will consist in taking the transition

leading to c' . This allows us to deduce that $c \in g(V_S^P)$. Assume now that $c \in C_P$. Then, because $c \in V_S^P$ and S is deadlock free, there necessarily exists c' such $c \rightarrow c'$ and $c' \in V_S^P$ (otherwise there would be a strategy to reach q_F from all states c' such that $c \rightarrow c'$ and hence c would not be in V_S^P). So also in this case we have $c \in g(V_S^P)$. Hence we have $V_S^P \subseteq g(V_S^P)$ which allows to deduce that $V_S^P \subseteq U$. \square

Lemma 9. $W_S^P = \llbracket \mu Y. \nu X. ((\bigvee_{q \in Q \setminus \{q_F\}} q) \wedge (Q_1 \vee \Diamond X) \wedge (Q_P \vee (Q_1 \wedge \Box X)) \vee (q_F \wedge Q_P \wedge \Diamond Y) \vee (q_F \wedge Q_1 \wedge \Box Y)) \rrbracket$.

Proof. We denote by U the set $\llbracket \mu Y. \nu X. ((\bigvee_{q \in Q \setminus \{q_F\}} q) \wedge (Q_1 \vee \Diamond X) \wedge (Q_P \vee (Q_1 \wedge \Box X)) \vee (q_F \wedge Q_P \wedge \Diamond Y) \vee (q_F \wedge Q_1 \wedge \Box Y)) \rrbracket$ and we consider the function $h : 2^C \mapsto 2^C$ such that for each set of configurations $C' \subseteq C$, we have $h(C') = \llbracket \nu X. ((\bigvee_{q \in Q \setminus \{q_F\}} q) \wedge (Q_1 \vee \Diamond X) \wedge (Q_P \vee (Q_1 \wedge \Box X)) \vee (q_F \wedge Q_P \wedge \Diamond Y) \vee (q_F \wedge Q_1 \wedge \Box Y)) \rrbracket_{\varepsilon_0[Y:=C']}$, i.e. it corresponds to the set $\llbracket \nu X. ((\bigvee_{q \in Q \setminus \{q_F\}} q) \wedge (Q_1 \vee \Diamond X) \wedge (Q_P \vee (Q_1 \wedge \Box X)) \vee (q_F \wedge Q_P \wedge \Diamond Y) \vee q_F \wedge (Q_1 \wedge \Box Y)) \rrbracket$ where Y is interpreted as C' . Note that U is then the least fixpoint of h .

We first prove that $U \subseteq W_S^P$. Note that U being the least fixpoint of h , by the Knaster-Tarski Theorem we have $U = \bigcap \{C' \subseteq C \mid h(C') \subseteq C'\}$. We will hence show that $h(W_S^P) \subseteq W_S^P$ from which we will get $U \subseteq W_S^P$. For this, we consider the function $g_{W_S^P} : 2^C \mapsto 2^C$ such that for each set of configurations $C' \subseteq C$, we have $g_{W_S^P}(C') = \llbracket (\bigvee_{q \in Q \setminus \{q_F\}} q) \wedge (Q_1 \vee \Diamond X) \wedge (Q_P \vee (Q_1 \wedge \Box X)) \vee (q_F \wedge Q_P \wedge \Diamond Y) \vee (q_F \wedge Q_1 \wedge \Box Y) \rrbracket_{\varepsilon_0[Y:=W_S^P, X:=C']}$. The set $V = h(W_S^P)$ is then by definition the greatest fixpoint of $g_{W_S^P}$ (so we have as well $V = g_{W_S^P}(V)$). Hence we need to show that $V \subseteq W_S^P$. For this we will assume that $c \notin W_S^P$ and show that $c \notin V$.

Let $c \notin W_S^P$. Hence there exists a strategy $\sigma \in \Sigma$ such that $\text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Box \Diamond q_F \rrbracket$. We reason now by contradiction assuming that $c \in V$. We will show that either there exists a finite play $c_0 \rightarrow c_1 \rightarrow c_2 \dots \rightarrow c_k$ in M_S which respects σ and with $c_k \in W_S^P \cap \llbracket q_F \rrbracket$ or there exists an infinite play $c_0 \rightarrow c_1 \rightarrow c_2 \rightarrow \dots$ in $\text{Plays}(M_S, c, \sigma)$ such that $c_0 = c$ and $c_i \notin \llbracket q_F \rrbracket$ for all $i \in \mathbb{N} \setminus \{0\}$.

We prove in fact by induction that if $c_0 \rightarrow c_1 \dots \rightarrow c_k$ is a finite play in M_S respecting σ with $c_0 = c$ and such that $c_i \in V$ and $c_i \notin W_S^P$ for $i \in [0..k]$, then there exists $c_{k+1} \in C$ such that either $c_{k+1} \in W_S^P \cap \llbracket q_F \rrbracket$ or ($c_{k+1} \in V$ and $c_{k+1} \notin \llbracket q_F \rrbracket$), and such that $c_0 \rightarrow c_1 \dots \rightarrow c_{k+1}$ is a play in M_S respecting σ . We proceed with the base case. First note that $c \in V$ and $c \notin W_S^P$. We recall that $V = g_{W_S^P}(V)$.

- If $c \in C_1$, then let $c_1 = \sigma(c)$. Since $c \in g_{W_S^P}(V)$ and $c \in C_1$, we have necessarily that either ($c_1 \in V$ and $c_1 \notin \llbracket q_F \rrbracket$) or $c_1 \in W_S^P \cap \llbracket q_F \rrbracket$, and we have $c \rightarrow c_1$ is a play in M_S respecting σ .
- If $c \in C_P$, then by definition of $g_{W_S^P}(V)$, there exists necessarily c_1 such that either ($c_1 \in V$ and $c_1 \notin \llbracket q_F \rrbracket$) or $c_1 \in W_S^P \cap \llbracket q_F \rrbracket$ and $c \rightarrow c_1$. Furthermore this allows us to deduce that $c \rightarrow c_1$ is a play in M_S respecting σ .

For the inductive case, the proof works exactly the same way.

From this we deduce that either there exists a finite play $c \rightarrow c_1 \rightarrow c_2 \dots \rightarrow c_k$ in M_S respecting σ with $c_k \in W_S^P \cap \llbracket q_F \rrbracket$ or there exists an infinite play $c \rightarrow c_1 \rightarrow c_2 \rightarrow \dots$ in $\text{Plays}(M_S, c, \sigma)$ such that $c_0 = c$ and $c_i \notin \llbracket q_F \rrbracket$ for all $i \in \mathbb{N}$. We recall that we have $\text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Box \Diamond q_F \rrbracket$ and proceed by a case analysis to show a contradiction:

1. If there exists $c_0 \rightarrow c_1 \rightarrow c_2 \dots \rightarrow c_k$ in $\text{Plays}(M_S, c, \sigma)$ with $c_k \in W_S^P$ then it is not possible that all the plays of the form $c_0 \rightarrow c_1 \rightarrow c_2 \dots \rightarrow c_k \dots$ in $\text{Plays}(M_S, c, \sigma)$ are in $\llbracket \Box \Diamond q_F \rrbracket$. In fact, otherwise there would be a strategy σ' from c_k (built from σ) such that $\text{Plays}(M_S, c_k, \sigma') \subseteq \llbracket \Box \Diamond q_F \rrbracket$ which contradicts the fact that $c_k \in W_S^P$.
2. If there exists an infinite play $c_0 \rightarrow c_1 \rightarrow c_2 \rightarrow \dots$ in $\text{Plays}(M_S, c, \sigma)$ such that $c_0 = c$ and $c_i \notin \llbracket q_F \rrbracket$ for all $i \in \mathbb{N}$, then we have a contradiction with the fact that $\text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Box \Diamond q_F \rrbracket$.

We hence deduce that $c \notin V$, and consequently we have shown that $U \subseteq W_S^P$.

We now prove that $W_S^P \subseteq U$. To do that we will instead show that the complement of U , denoted by \overline{U} , is included in the complement of W_S^P , denoted by W_S^1 and which is equal to the set $\{c \in C \mid \exists \sigma \in \Sigma \text{ s.t. } \text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Box \Diamond q_F \rrbracket\}$.

First, note that \overline{U} is equal to $\llbracket \nu Y. \mu X. (q_F \vee (Q_P \wedge \Box X) \vee (Q_1 \wedge (Q_P \vee \Diamond X)) \wedge (\neg q_F \vee Q_P \vee \Diamond Y) \wedge (\neg q_F \vee Q_1 \vee \Box Y)) \rrbracket$, which is itself equal to $\llbracket \nu Y. \mu X. [(\neg q_F \wedge ((Q_P \wedge \Box X) \vee (Q_1 \wedge \Diamond X))) \vee (q_F \wedge ((Q_P \wedge \Box Y) \vee (Q_1 \vee \Diamond Y)))] \rrbracket$. Note that then we have as well that $\overline{U} = \llbracket \mu X. [(\neg q_F \wedge ((Q_P \wedge \Box X) \vee (Q_1 \wedge \Diamond X))) \vee (q_F \wedge ((Q_P \wedge \Box Y) \vee (Q_1 \vee \Diamond Y)))] \rrbracket_{\varepsilon_0[Y:=\overline{U}]}$. We denote by T the set $\llbracket (q_F \wedge ((Q_P \wedge \Box Y) \vee (Q_1 \vee \Diamond Y))) \rrbracket_{\varepsilon_0[Y:=\overline{U}]}$. Note that we have : $\llbracket \mu X. [(\neg q_F \wedge ((Q_P \wedge \Box X) \vee (Q_1 \wedge \Diamond X))) \vee (q_F \wedge ((Q_P \wedge \Box Y) \vee (Q_1 \vee \Diamond Y)))] \rrbracket_{\varepsilon_0[Y:=\overline{U}]} \subseteq \llbracket \mu X. [(((Q_P \wedge \Box X) \vee (Q_1 \wedge \Diamond X))) \vee (q_F \wedge ((Q_P \wedge \Box Y) \vee (Q_1 \vee \Diamond Y)))] \rrbracket_{\varepsilon_0[Y:=\overline{U}]}$. Let U' be the set $\llbracket \mu X. [(((Q_P \wedge \Box X) \vee (Q_1 \wedge \Diamond X))) \vee (q_F \wedge ((Q_P \wedge \Box Y) \vee (Q_1 \vee \Diamond Y)))] \rrbracket_{\varepsilon_0[Y:=\overline{U}]}$. By adapting the proof of Lemma 8, we can deduce that the complement of this last U' is equal to $\{c \in C \mid \exists \sigma \in \Sigma \text{ s.t. } \text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Diamond T \rrbracket\}$ (where $\llbracket \Diamond T \rrbracket$ denotes the plays that eventually reach the set T) and consequently $U' = \{c \in C \mid \exists \sigma \in \Sigma \text{ s.t. } \text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Diamond T \rrbracket\}$. We have consequently that $\overline{U} \subseteq \{c \in C \mid \exists \sigma \in \Sigma \text{ s.t. } \text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Diamond T \rrbracket\}$ with $T = \llbracket (q_F \wedge ((Q_P \wedge \Box Y) \vee (Q_1 \vee \Diamond Y))) \rrbracket_{\varepsilon_0[Y:=\overline{U}]}$.

We can now prove that $\overline{U} \subseteq W_S^1$. Let $c \in \overline{U}$. Since $\overline{U} \subseteq \{c \in C \mid \exists \sigma \in \Sigma \text{ s.t. } \text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Diamond T \rrbracket\}$, from c Player 1 can surely reach T and by definition of T , when it reaches T first it is in $\llbracket q_F \rrbracket$ and then Player 1 can ensure a successor state to belong to \overline{U} . Hence performing this reasoning iteratively, we can build a strategy for Player 1 from c to reach surely infinitely often $\llbracket q_F \rrbracket$. \square

B Proofs of Section 4

B.1 Proof of Lemma 2

Proof. Let $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ be a 1-VASS-MDP and $M_S = \langle C, C_1, C_P, \rightarrow, p \rangle$ its associated MDP. If a configuration $\langle q, \mathbf{v} \rangle \in C_P$ is a deadlock, then it means that there is no outgoing edge in S from the control state q ; hence each time a play will reach this configuration, Player 1 will lose, so we can add a self loop without any effect on the counters to this state in order to remove the deadlock. For the states $q \in Q_1$, we add a transition t to the outgoing edge of q which does not modify the counter values and which leads to a new control state with a self-loop, such that if the play reaches a configuration $\langle q, \mathbf{v} \rangle$ which is a deadlock in S , in the new game arena the only choice for Player 1 will be to go to this new absorbing state and he will lose as he loses in S because of the deadlock. \square

B.2 Proof of Theorem 6

If we define the two following set of configurations: $V_S^1 = \{c \in C \mid \exists \sigma \in \Sigma \text{ such that } \text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Diamond q_F \rrbracket\}$ and $W_S^1 = \{c \in C \mid \exists \sigma \in \Sigma \text{ such that } \text{Plays}(M_S, c, \sigma) \subseteq \llbracket \Box \Diamond q_F \rrbracket\}$, we have the following result:

Lemma 10.

- $V_S^1 = \llbracket \mu X. q_F \vee (Q_1 \wedge \Diamond X) \vee (Q_P \wedge \Box X) \rrbracket$
- $W_S^1 = \llbracket \nu Y. \mu X. (q_F \vee (Q_P \wedge \Box X) \vee (Q_1 \wedge \Diamond X)) \wedge (\bigvee_{q \in Q \setminus \{q_F\}} q \vee Q_1 \vee (Q_P \wedge \Box Y)) \wedge (\bigvee_{q \in Q \setminus \{q_F\}} q \vee Q_P \vee \Diamond Y) \rrbracket$

Proof. Since we are looking at deadlock free VASS-MDP, we can reuse the formula given by Lemma 8 and 9. In fact, note that we have $V_S^1 = \overline{V_S^P}$ and $W_S^1 = \overline{W_S^P}$. Hence by taking the complement formulae of L_μ^{sv} , we obtain the desired result. For V_S^1 , from the formula describing V_S^P , we obtain the formula $\mu X. q_F \vee (Q_1 \wedge \Diamond X) \vee (Q_1 \wedge (Q_P \vee \Box X))$ which is equivalent to $\mu X. q_F \vee (Q_1 \wedge \Diamond X) \vee (Q_P \wedge \Box X)$. For W_S^1 , from the formula describing W_S^P , we obtain the formula $\nu Y. \mu X. (q_F \vee (Q_P \wedge \Box X) \vee (Q_1 \wedge (Q_P \vee \Diamond X))) \wedge (\bigvee_{q \in Q \setminus \{q_F\}} q \vee Q_1 \vee \Box Y) \wedge (\bigvee_{q \in Q \setminus \{q_F\}} q \vee Q_P \vee \Diamond Y)$ which is equivalent to $\nu Y. \mu X. (q_F \vee (Q_P \wedge \Box X) \vee (Q_1 \wedge \Diamond X)) \wedge (\bigvee_{q \in Q \setminus \{q_F\}} q \vee Q_1 \vee (Q_P \wedge \Box Y)) \wedge (\bigvee_{q \in Q \setminus \{q_F\}} q \vee Q_P \vee \Diamond Y)$. \square

B.3 Closed L_μ^{sv} -formulae manipulate upward closed-sets

We will say that an environment $\varepsilon : \mathcal{X} \rightarrow 2^C$ is upward-closed if for each variable $X \in \mathcal{X}$, $\varepsilon(X)$ is upward-closed (we take as order \leq for the configurations, the classical one such that $\langle q, \mathbf{v} \rangle \leq \langle q', \mathbf{v}' \rangle$ iff $q = q'$ and $\mathbf{v} \leq \mathbf{v}'$). Whereas it is not true that for any VASS, any formula $\phi \in L_\mu^{sv}$ and any upward closed environment $\varepsilon : \mathcal{X} \rightarrow 2^C$ the set $\llbracket \phi \rrbracket_\varepsilon$ is upward closed, we now prove that on single-sided VASS this property holds.

Lemma 11. *For any formula $\phi \in L_\mu^{sv}$ and any upward closed environment ε , $\llbracket \phi \rrbracket_\varepsilon$ evaluated over the configurations of the (Q_1, Q_P) -single-sided VASS $\langle Q, T \rangle$ is an upward closed set.*

Proof. The proof is by induction on the length of the formula ϕ . For formulae of the form q , the result is due to the fact that the set of considered regions is upward closed. For formulae of the form X , the result comes from the assumption on the considered environment. For formulae of the form $\phi \wedge \psi$ and $\mu X.\phi$ the result can be obtained using the induction hypothesis and the fact that the intersection of upward closed sets is an upward-closed set. For formulae of the form $\phi \vee \psi$ and $\nu X.\phi$ the result can be obtained using the induction hypothesis and the fact that the union of upward closed sets is an upward closed set.

Now we consider formulae of the form $\Diamond\phi$ assuming that for any upward closed environment ε , $\llbracket \phi \rrbracket_\varepsilon$ is an upward closed set. Let $c_1 \in \llbracket \Diamond\phi \rrbracket_\varepsilon$. Then there exists $c'_1 \in \llbracket \phi \rrbracket_\varepsilon$ such that $c_1 \rightarrow c'_1$. Let $c_2 \in C$ such that $c_1 \leq c_2$. Since we are considering VASS, we now that there exists $c'_2 \in C$ such that $c_2 \rightarrow c'_2$ and $c'_1 \leq c'_2$. Since $\llbracket \phi \rrbracket_\varepsilon$ is upward closed, we have $c'_2 \in \llbracket \phi \rrbracket_\varepsilon$, hence c_2 belongs to $\llbracket \Diamond\phi \rrbracket_\varepsilon$. This proves that $\llbracket \Diamond\phi \rrbracket_\varepsilon$ is upward closed.

Now we consider formulae of the form $Q_P \wedge \Box\phi$ assuming that for any upward closed environment ε , $\llbracket \phi \rrbracket_\varepsilon$ is an upward closed set. Let $c_1 \in \llbracket Q_P \wedge \Box\phi \rrbracket_\varepsilon$. Then for all $c'_1 \in C$ such that $c_1 \rightarrow c'_1$, we have $c'_1 \in \llbracket \phi \rrbracket_\varepsilon$. Note that $c_1 = (q_1, \mathbf{v}_1)$ with $q_1 \in Q_P$. Let $c_2 \in C$ such that $c_1 \leq c_2$. By definition of the order \leq on the set of configurations, we have $c_2 = (q_1, \mathbf{v}_2)$ with $\mathbf{v}_1 \leq \mathbf{v}_2$. Let c'_2 such that $c_2 \rightarrow c'_2$. Since $q_1 \in Q_P$, by the definition of single-sided VASS, we know that the transition which leads from c_2 to c'_2 can also be taken from c_1 (this is because the outgoing transitions from control states in Q_P do not modify the counter values), hence there exists $c'_1 \in C$ such that $c_1 \rightarrow c'_1$ and since $c_1 \leq c_2$, we have $c'_1 \leq c'_2$. Furthermore, we have $c'_1 \in \llbracket \phi \rrbracket_\varepsilon$ and since, by induction, this last set is upward closed, we deduce $c'_2 \in \llbracket \phi \rrbracket_\varepsilon$. This allows us to conclude that c_2 belongs to $\llbracket Q_1 \wedge \Box\phi \rrbracket_\varepsilon$ which is hence an upward closed set. \square

B.4 Proof of Lemma 3

Proof. We consider the function $h : 2^C \rightarrow 2^C$ which associates to each set of configurations $C' \subseteq C$ the set $h(C') = \llbracket q_F \vee \text{InvPre}(X, Y) \rrbracket_{\varepsilon_0[X:=E, Y:=C']}$. We define a sequence of sets $(F_i)_{i \in \mathbb{N}}$ included in C as follows:

- $F_0 = \emptyset$
- for all $i \in \mathbb{N}$, $F_{i+1} = F_i \cup h(F_i)$

Using Lemma 11 and the fact that the union of upward closed set is an upward closed set, we can prove that F_i is upward closed for all $i \in \mathbb{N}$. Furthermore, we have that $F_i \subseteq F_{i+1}$ for all $i \in \mathbb{N}$. Since $(F_i)_{i \in \mathbb{N}}$ is an increasing sequence of upward closed sets included in C and since (C, \leq) is a wqo, from the theory of wqo, we know that there exists $N \in \mathbb{N}$ such that for all $i \geq N$, $F_i = F_{i+1}$.

We consider also the function $g : 2^C \rightarrow 2^C$, which associates to each set of configurations $C' \subseteq C$ the set $g(C') = \llbracket \mu Y.(q_F \vee \text{InvPre}(X, Y)) \rrbracket_{\varepsilon_0[X:=C']}$. By

definition, E is the greatest fixpoint of the function g , hence $E = \llbracket \mu Y. (q_F \vee \text{InvPre}(X, Y)) \rrbracket_{\varepsilon_0[X:=E]}$. E is then also the least fixpoint of the function h and consequently, by definition of the sequence $(F_i)_{i \in \mathbb{N}}$, we know that $E = \bigcup_{i \in \mathbb{N}} F_i$. This allows us to deduce that $E = F_N$. We point out that $F_1 = \llbracket q_F \rrbracket$.

We now define a strategy σ for Player 1 which will be memoryless on E (i.e. the strategy will only depend on the current configuration) and is a function $\sigma : E \cap C_1 \mapsto C$ (note that since with this strategy all the plays starting from E will stay in E , we do not need to define it precisely on the entire set C_1 and we assume that, on the set $C_1 \setminus E$, the strategy can choose any one of the possible successor configurations). Let $c \in E$. Then we denote by $j \in [2..N]$ the smallest index such that $c \in F_j$ and $c \notin F_{j-1}$, we then define $\sigma(c)$ as being the configuration c' such that $c \rightarrow c'$ and $c' \in F_{j-1}$ (by definition of the sequence $((F_i)_{i \in \mathbb{N}})$, such a c' necessarily exists). If c belongs to F_1 , then the strategy chooses any one of the possible successors.

Let $c_0 \in E$. We show that there exists a play $c_0 \cdot c_1 \cdot c_2 \dots$ in $\text{Plays}(M_S, c_0, \sigma)$ that satisfies the three properties of the lemma. In fact, we consider the play such that in all configurations $c \in E \cap C_P$ and if $j \in [2..N]$ is the smallest index such that $c \in F_j$ and $c \notin F_{j-1}$, Player P chooses c' such that $c \rightarrow c'$ and $c' \in F_{j-1}$. Hence in this play it is obvious that in less than N steps, the play will reach a configuration in $F_1 = \llbracket q_F \rrbracket$ and the points 2. and 3. also hold for this play by definition of the sets $(F_i)_{i \in \mathbb{N}}$. \square

B.5 Proof of Lemma 4

We now prove that the set E is included in V_{AS}^1 . Techniques we used here are quite similar than the one presented in [2] to prove decidability of the sure reachability problems in probabilistic VASS (without nondeterminism).

Lemma 12. $E \subseteq V_{AS}^1$

Proof. We consider the integer $N \in \mathbb{N}$ and the strategy σ of Player 1 given by Lemma 3. For each $q \in Q_P$, let $\text{Out}(q) = \{(q, \mathbf{z}, q') \in T \text{ for some } q' \in Q \text{ and } \mathbf{z} \in \mathbb{Z}^n\}$ be the set of transitions going out of q . We also denote by L_q the cardinality of $\text{Out}(q)$, by W_q the sums $\sum_{t \in \text{Out}(q)} \tau(t)$ and finally Min_q is the minimal element of $\{\tau(t) \mid t \in \text{Out}(q)\}$. By definition of VASS-MDP, we know that for a configuration $\langle q, \mathbf{v} \rangle \in C_P$, for any configuration $c' \in C$ such that $\langle q, \mathbf{v} \rangle \rightarrow c'$, we have $p(q, \mathbf{v})(c') \geq \frac{\text{Min}_q}{L_q \cdot W_q}$. We denote by β the minimal element of the set $\{\frac{\text{Min}_q}{L_q \cdot W_q} \mid q \in Q_P\}$. Then for any configuration $c \in C_P$ and $c' \in C$ such that $c \rightarrow c'$, we have $p(c)(c') \geq \beta$. Note that necessarily $\beta > 0$. Let $c_0 \in E \setminus \llbracket q_F \rrbracket$ and let $c_0 \cdot c_1 \cdot c_2 \dots$ be a play in $\text{Plays}(M_S, c_0, \sigma)$ such that for all $i \in \mathbb{N}$, $c_i \notin \llbracket q_F \rrbracket$. From Lemma 3, we know that there exists $N \in \mathbb{N}$ such that, for all $i \in \mathbb{N}$, $c_i \in E$, we have that $\mathbb{P}(M_S, c_i, \sigma, \llbracket \Diamond q_F \rrbracket) \geq \beta^N$. This allows us to deduce that the probability of never visiting q_F from c_0 following σ is smaller than $(1 - \beta^N)^\infty$ and since $\beta > 0$, we deduce that this probability is equal to 0. Consequently $\mathbb{P}(M_S, c_0, \sigma, \llbracket \Diamond q_F \rrbracket) = 1$ and $c_0 \in V_{AS}^1$. \square

We now prove the opposite direction. For this we use a technique similar to the ones presented in the proof of Lemma 5.29 in [5].

Lemma 13. $V_{AS}^1 \subseteq E$

Proof. Let $c_0 \in V_{AS}^1$. So there exists a strategy σ for Player 1 such that $\mathbb{P}(M_S, c_0, \sigma, \llbracket \Diamond q_F \rrbracket) = 1$. Let D be the following set of configurations: $\{c \in C \mid \exists c_0 \cdot c_1 \cdots \in \text{Plays}(M_S, c_0, \sigma) \text{ s.t. } \exists i \in \mathbb{N} \text{ for which } c_i = c \text{ and } \forall 0 \leq j < i. c_j \notin \llbracket q_F \rrbracket\}$. Clearly c_0 belongs to D . We will show that $D \subseteq E$.

We consider the two functions $g, h : 2^C \rightarrow 2^C$ such that for each set of configurations $C' \subseteq C$, we have $g(C') = \llbracket \mu Y. (q_F \vee \text{InvPre}(X, Y)) \rrbracket_{\varepsilon_0[X:=C']}$ and $h(C') = \llbracket q_F \vee \text{InvPre}(X, Y) \rrbracket_{\varepsilon_0[X:=D, Y:=C']}$. We define the following sequence of configurations $(F_i^D)_{i \in \mathbb{N}}$ such that:

- $F_0^D = \emptyset$,
- for all $i \in \mathbb{N}$, $F_{i+1}^D = F_i^D \cup h(F_i^D)$

Let $c \in D$, since $\mathbb{P}(M_S, c_0, \sigma, \llbracket \Diamond q_F \rrbracket) = 1$, we know that there exists a strategy σ' such that $\mathbb{P}(M_S, c, \sigma', \llbracket \Diamond q_F \rrbracket) = 1$ (otherwise we would have $\mathbb{P}(M_S, c_0, \sigma, \llbracket \Diamond q_F \rrbracket) < 1$). Hence there is a play in $c \cdot c'_1 \cdot c'_2 \cdots \in \text{Plays}(M_S, c, \sigma')$ for which there exists $i \in \mathbb{N}$ satisfying $c'_i \in \llbracket q_F \rrbracket$ and $c'_j \notin \llbracket q_F \rrbracket$ for all $0 \leq j < i$. Note also that by definition of D , for all $0 \leq j \leq i$, we have that c'_j belongs to D and if $c'_j \in C_P$ then for all $c'' \in C$ such that $c'_j \rightarrow c''$, we have $c'' \in D$. Hence there exists $0 \leq k \leq i + 1$, such that $c \in F_k^D$. Since $g(D) = \llbracket \mu Y. (q_F \vee \text{InvPre}(X, Y)) \rrbracket_{\varepsilon_0[X:=D]}$, we have $g(D) = \bigcup_{i \in \mathbb{N}} F_i^D$. Consequently, $c \in g(D)$ and we have $D \subseteq g(D)$.

We know that E is the greatest fixpoint of g , then by Knaster-Tarski Theorem we know that $E = \bigcup \{C' \in C \mid C' \subseteq g(C')\}$, and hence $D \subseteq E$. Since $c_0 \in D$, we deduce that $c_0 \in E$. \square

B.6 Proof of Lemma 5

We denote by F the set $\llbracket \nu X. \text{InvPre}(X, \mu Y. (q_F \vee \text{InvPre}(X, Y))) \rrbracket_{\varepsilon_0}$. Since $\nu X. \text{InvPre}(X, \mu Y. (q_F \vee \text{InvPre}(X, Y)))$ is a formula of L_μ^{sv} interpreted over the (Q_1, Q_P) -single-sided VASS $\langle Q, T \rangle$, from Lemma 11, we know that F is an upward-closed set.

We first prove that W_{AS}^1 is included in F .

Lemma 14. $W_{AS}^1 \subseteq F$

Proof. Let $c_0 \in W_{AS}^1$. So there exists a strategy σ for Player 1 such that $\mathbb{P}(M_S, c_0, \sigma, \llbracket \Box \Diamond q_F \rrbracket) = 1$. Let T be the following set of configurations $\{c \in C \mid \exists c_0 \cdot c_1 \cdots \in \text{Plays}(M_S, c_0, \sigma) \text{ s.t. } \exists i \in \mathbb{N} \text{ for which } c_i = c\}$. Necessarily, we have $T \cap \llbracket q_F \rrbracket \neq \emptyset$ and $c_0 \in T$. We consider the function $g : 2^C \rightarrow 2^C$ such that for each set of configurations $C' \subseteq C$, we have $g(C') = \llbracket \text{InvPre}(X, \mu Y. (q_F \vee \text{InvPre}(X, Y))) \rrbracket_{\varepsilon_0[X:=C']}$. We will prove that $T \subseteq g(T)$.

Let $c \in T$. We have necessarily that there exists a strategy σ' for Player 1 such that $\mathbb{P}(M_S, c, \sigma', \llbracket \Box \Diamond q_F \rrbracket) > 0$, otherwise we would have that $\mathbb{P}(M_S, c_0, \sigma, \llbracket \Box \Diamond q_F \rrbracket) < 1$. Hence there exists a play in M_S respecting σ of the form:

$$c_0 \cdot c_1 \cdots c_k \cdots c \cdot c'_1 \cdots c'_m$$

with $m \geq 1$ (that is $c \neq c'_m$) and $c'_m \in \llbracket q_F \rrbracket$. By definition of T , for each $1 \leq i \leq m$, if $c'_i \in C_P$, for all configurations $c'' \in C$ such that $c'_i \rightarrow c''$ we have, $c'' \in T$. Using a similar reasoning to that done in Lemma 13, we deduce that for all $1 \leq i \leq m$, we have $c'_i \in \llbracket \mu Y.(q_F \vee \text{InvPre}(X, Y)) \rrbracket_{\varepsilon_0[X:=T]}$. Furthermore if $c \in C_P$, for all configurations $c'' \in C$ such that $c \rightarrow c''$ we have also $c'' \in T$, hence we have that $c \in g(T)$ (using the definition of $\text{InvPre}(X, Y)$). This implies $T \subseteq g(T)$.

Since F is the greatest fixpoint of g , then by Knaster-Tarski Theorem we know that $F = \bigcup \{C' \in C \mid C' \subseteq g(C')\}$, so $T \subseteq F$. Since $c_0 \in T$, we deduce that $c_0 \in F$. \square

We now prove the left to right inclusion.

Lemma 15. $F \subseteq W_{AS}^1$

Proof. We consider the function $h : 2^C \rightarrow 2^C$ which associates to each set of configurations $C' \subseteq C$ the set $h(C') = \llbracket q_F \vee \text{InvPre}(X, Y) \rrbracket_{\varepsilon_0[X:=F, Y:=C']}$. We define a sequence of sets $(F_i)_{i \in \mathbb{N}}$ included in C as follows:

- $F_0 = \emptyset$
- for all $i \in \mathbb{N}$, $F_{i+1} = F_i \cup h(F_i)$

As for the proof of Lemma 3, we know that there exists $N \in \mathbb{N}$ such that for all $i \geq N$, $F_i = F_{i+1}$ and that $\llbracket \mu Y.(q_F \vee \text{InvPre}(X, Y)) \rrbracket_{\varepsilon_0[X:=F]} = F_N$. Since $F = \llbracket \nu X.\text{InvPre}(X, \mu Y.(q_F \vee \text{InvPre}(X, Y))) \rrbracket_{\varepsilon_0}$, using again fixpoint theory, we know that $F = \llbracket \text{InvPre}(X, \mu Y.(q_F \vee \text{InvPre}(X, Y))) \rrbracket_{\varepsilon_0[X:=F]}$ and consequently $F = \llbracket \text{InvPre}(X, Y) \rrbracket_{\varepsilon_0[X:=F, Y:=F_N]}$.

We now define a strategy σ for player 0 which will be memoryless on F (i.e. the strategy will only depend of the current configuration). The strategy σ will be described as a function $\sigma : F \cap C_1 \mapsto C$ (note that since with this strategy all the plays starting from F will stay in F , we do not need to define it precisely on the entire set C_1 ; we assume that on the set $C_1 \setminus F$, the strategy can choose any one of the possible successor states). Let $c \in C_1 \setminus F$ and consider the following two cases. If $c \in F \setminus F_N$ or $c \in F_1$, we define $\sigma(c)$ as being a configuration $c' \in F_N \cap F$ such that $c \rightarrow c'$. By definition of F such a configuration necessarily exists. If $c \in F_N \setminus F_1$ we denote by $j \in [2..N]$ the smallest index such that $c \in F_j$ and $c \notin F_{j-1}$. We then define $\sigma(c)$ as being the configuration c' such that $c \rightarrow c'$ and $c' \in F_{j-1} \cap F$ (by definition of F and of the sequence $(F_i)_{i \in \mathbb{N}}$, such a c' necessarily exists).

By construction of the strategy σ and using a similar reasoning to the one performed in the proof of Lemma 12, we can prove that for all $c \in F$, we have $\mathbb{P}(M_S, c, \sigma, \llbracket \Diamond q_F \rrbracket) = 1$. We would like now to prove that for each $c \in F$, $\mathbb{P}(M_S, c, \sigma, \llbracket \Box \Diamond q_F \rrbracket) = 1$.

We denote by $\llbracket \Box \neg q_F \rrbracket$ the set of infinite plays $c_0 \cdot c_1 \cdots$ of M_S such that for all $i \in \mathbb{N}$, $c_i \notin \llbracket q_F \rrbracket$. Then for each $c \in F$, we have $\mathbb{P}(M_S, c, \sigma, \llbracket \Box \neg q_F \rrbracket) = 0$. We also use the notation $\llbracket \Diamond \Box \neg q_F \rrbracket$ the set of infinite play $c_0 \cdot c_1 \cdots$ of M_S for which there exists $i \in \mathbb{N}$, such that $c_j \notin \llbracket q_F \rrbracket$ for all $j \geq i$. Then for each $c \in F$, we

have $\mathbb{P}(M_S, c, \sigma, \llbracket \Diamond q_F \rrbracket) = 1 - \mathbb{P}(M_S, c, \sigma, \llbracket \Diamond \neg q_F \rrbracket)$. We will now prove that for each $c \in F$, $\mathbb{P}(M_S, c, \sigma, \llbracket \Diamond \neg q_F \rrbracket) = 0$.

For $c \in C$, $i \in \mathbb{N}$ and $d \in C \setminus \llbracket q_F \rrbracket$, let Π_{c-i-d} be the set of finite plays of M_S of the form $c_0 \cdot c_1 \cdots c_k$ such that:

- $k > i$;
- $c_0 = c$, $c_k = d$ and $c_{k-1} \in \llbracket q_F \rrbracket$;
- the set $G = \{j \in \{0, \dots, k\} \mid c_j \in \llbracket q_F \rrbracket\}$ has i elements.

This represents the set of finite plays starting at c ending at configuration d and that passes exactly through i configurations in $\llbracket q_F \rrbracket$.

For $c \in F$ and $i \in \mathbb{N}$ and $d \in F \setminus \llbracket q_F \rrbracket$, we define Δ_{c-i-d} the set of infinite plays of the form $\rho \cdot \rho'$ where ρ is a finite play in Π_{c-i-d} and ρ' is an infinite play in $\llbracket \Box \neg q_F \rrbracket$. Let $\Delta_{c-i} = \bigcup_{d \in F \setminus \llbracket q_F \rrbracket} \Delta_{c-i-d}$. Intuitively Δ_{c-i} is the set of infinite plays starting from c which revisits $\llbracket q_F \rrbracket$ exactly i times. For $c \in F$, it is straightforward to check that:

- (1) $\forall i \in \mathbb{N}, \forall d_1, d_2 \in F \setminus \llbracket q_F \rrbracket$ such that $d_1 \neq d_2$, $\Delta_{c-i-d_1} \cap \Delta_{c-i-d_2} = \emptyset$
- (2) $\forall i, j \in \mathbb{N}$ such that $i \neq j$, $\Delta_{c-i} \cap \Delta_{c-j} = \emptyset$
- (3) $\forall i \in \mathbb{N}, \forall d \in F \setminus \llbracket q_F \rrbracket$, $\mathbb{P}(M_S, c, \sigma, \Delta_{c-i-d}) = P(\Pi_{c-i-d})\mathbb{P}(M_S, d, \sigma, \llbracket \Box \neg q_F \rrbracket)$

Hence, for all $c \in F$ and $i \in \mathbb{N}$, we have:

$$\begin{aligned} \mathbb{P}(M_S, c, \sigma, \Delta_{c-i}) &= \sum_{d \in F \setminus \llbracket q_F \rrbracket} \mathbb{P}(M_S, c, \sigma, \Delta_{c-i-d}) \\ &= \sum_{d \in F \setminus \llbracket q_F \rrbracket} P(\Pi_{c-i-d})\mathbb{P}(M_S, d, \sigma, \llbracket \Box \neg q_F \rrbracket) \\ &= 0 \end{aligned}$$

where the first equality holds by (1) and by the fact that all the configurations reached from c following σ belongs to F (by definition of σ and F); the second equality follows from (3) and the last equality from the fact that for all $d \in F$, $\mathbb{P}(M_S, d, \sigma, \llbracket \Box \neg q_F \rrbracket) = 0$.

Finally, we have for all $c \in F$:

$$(\text{Plays}(M_S, c, \sigma) \cap \llbracket \Diamond \Box \neg q_F \rrbracket) \subseteq (\text{Plays}(M_S, c, \sigma) \cap \bigcup_{i \in \mathbb{N}} \Delta_{c-i})$$

This is due to the fact that if an infinite play belongs to $\text{Plays}(M_S, c, \sigma) \cap \llbracket \Diamond \Box \neg q_F \rrbracket$, then it will pass only a finite number of times through $\llbracket q_F \rrbracket$. From this inclusion, the previous equality and using (2), we deduce that $\mathbb{P}(M_S, c, \sigma, \llbracket \Diamond \Box \neg q_F \rrbracket) \leq \sum_{i \in \mathbb{N}} \mathbb{P}(M_S, c, \sigma, \Delta_{c-i}) = 0$. Hence, for all $c \in F$, we have $\mathbb{P}(M_S, c, \sigma, \llbracket \Diamond q_F \rrbracket) = 1$, which allows us to conclude that $F \subseteq W_{AS}^1$. \square

B.7 Proof of Lemma 6

Proof. Algorithm 1 explores an unfolding of the computation tree of S , which is finitely branching since $|T|$ is finite. The number of counters is fixed, and therefore, by Dickson's Lemma, (\mathbb{N}^d, \preceq) is a well quasi ordering. Therefore, on every branch we eventually satisfy either the condition of line 19 or of line 7. In

the former case, a loop in the derived system S' is created, and the exploration of the current branch stops. In the latter case, a finitary description of a new colored (possibly infinite-state) subsystem is added to S' by adding finitely many states, transitions and configurations to Q' , T' and X' , respectively. Also in this case, the exploration of the current branch stops. Since the exploration is finitely branching, and every branch eventually stops, the algorithm terminates. \square

B.8 Proof of Lemma 7

Lemma 16. $\mathbb{P}^+(M_S, c_0, \llbracket \Diamond X \rrbracket) = 1 \implies \mathbb{P}^+(M_{S'}, c'_0, \llbracket \Diamond X' \rrbracket) = 1.$

Proof. Let us assume that $\mathbb{P}^+(M_S, c_0, \llbracket \Diamond X \rrbracket) = 1$. Therefore, there exists a family of strategies that make the probability of reaching X arbitrarily close to 1. In other words, $\forall \epsilon, \exists \sigma_\epsilon, \mathbb{P}(M_S, c_0, \sigma_\epsilon, \llbracket \Diamond X \rrbracket) \geq 1 - \epsilon$. For every $\epsilon > 0$ we use the strategy σ_ϵ of player 1 on M_S to construct a copycat strategy σ'_ϵ for the game on $M_{S'}$ that starts in $c'_0 = (q_0, \mathbf{0})$, such that it achieves $\llbracket \Diamond X' \rrbracket$ with probability $\geq 1 - \epsilon$.

The strategy σ'_ϵ will use the same moves on $M_{S'}$ as σ_ϵ on M_S , which is possible due to the way how $M_{S'}$ is constructed from M_S by Algorithm 1. By construction, for every reachable configuration in M_S there is a corresponding configuration in $M_{S'}$, and this correspondence can be maintained stepwise in the moves of the game.

For the initial uncolored part of $M_{S'}$, this is immediate, since S' is derived from the unfolding of the game tree of S . The correspondence is expressed by the function λ . Each current state of $M_{S'}$ is labeled by the corresponding current configuration of M_S .

In the colored subsystems, the corresponding configuration in system $M_{S'}$ is a projection of a configuration in M_S . For any transition $t \in T$ that is controlled by player 1 from a configuration in M_S , there exists a transition $t' \in T'$ that belongs to player 1 in the corresponding configuration in $M_{S'}$, such that this transition leads to the corresponding state. This is achieved by the projection and the fact that the 1-VASS-MDP game is monotone w.r.t. player 1, i.e., larger configurations always benefit the player (by allowing the same moves or even additional moves).

We now show a property on how probabilistic transitions in M_S and $M_{S'}$ correspond to each other: For every probabilistic transition $t \in T$ from a configuration in M_S , there exists a probabilistic transition $t' \in T'$ in the corresponding configuration in $M_{S'}$, and vice-versa, such that these transitions have the same probability. In particular, a configuration in $M_{S'}$ does not allow any additional probabilistic transitions compared to its corresponding configuration in M_S (though it may allow additional transitions controlled by player 1).

The first part of this statement follows from the monotonicity of the projection function and the monotonicity of the transitions w.r.t. the size of the configurations. For the second part we need to show that for every probabilistic transition $t' = \langle col_i(x), proj_k(op), col_i(y) \rangle \in T'$ from a configuration in $M_{S'}$,

there exists a probabilistic transition $t = \langle x, \mathbf{0}, y \rangle \in T_P$ in the corresponding configuration in S , such that the probabilities of these transitions are equal. This latter fact holds only because we are considering 1-VASS-MDP, where only the player can change the counters, whereas the probabilistic transitions can only change the control-states. I.e., the ‘larger’ projected configurations in $M_{S'}$ do not enable additional probabilistic transitions, since in 1-VASS-MDP these only depend on the control-state.

Therefore, by playing in $M_{S'}$ using strategy σ'_ϵ with the same moves as σ_ϵ plays in M_S , we reach the same corresponding configurations in $M_{S'}$ with the same probability values as in M_S . Since the definition of the target set X' in S' includes all configurations corresponding to configurations in X on S , it follows from $\mathbb{P}(M_S, c_0, \sigma_\epsilon, \llbracket \Diamond X \rrbracket) \geq 1 - \epsilon$ that $\mathbb{P}(M_{S'}, c'_0, \sigma'_\epsilon, \llbracket \Diamond X' \rrbracket) \geq 1 - \epsilon$. Since, by assumption above, this holds for every $\epsilon > 0$, we obtain $\mathbb{P}^+(M_{S'}, c'_0, \llbracket \Diamond X' \rrbracket) = 1$. \square

Lemma 17. $\mathbb{P}^+(M_{S'}, c'_0, \llbracket \Diamond X' \rrbracket) = 1 \implies \mathbb{P}^+(M_S, c_0, \llbracket \Diamond X \rrbracket) = 1$.

Proof. We use the assumed family of strategies on $M_{S'}$ that witnesses the property $\mathbb{P}^+(M_{S'}, c'_0, \llbracket \Diamond X' \rrbracket) = 1$ to synthesize a family of strategies on M_S that witnesses $\mathbb{P}^+(M_S, c_0, \llbracket \Diamond X \rrbracket) = 1$.

First we establish some basic properties of the system S' . It is a 1-VASS-MDP of dimension $d - 1$ with initial configuration c'_0 , and consists of several parts. The initial uncolored part induces a finite-state MDP. Moreover, S' contains finitely many subsystems of distinct colors, where each subsystem is a 1-VASS-MDP of dimension $d - 1$ obtained from S by projecting out one component of the integer vector. For color i , let $k(i)$ be the projected component of the vector (see line 9 of the algorithm). Each colored subsystem of dimension $d - 1$ induces an MDP that may be infinite-state (unless $d = 1$, in which case it is finite-state).

Note that colored subsystems are not reachable from each other, i.e., a color, once reached, is preserved. Each colored subsystem has its own initial configuration (created in lines 11-12 of Alg. 1). Let m be the number of colors in S' and r_i the initial configuration of the subsystem of color i (where $0 \leq i \leq m - 1$).

Let's now consider only those colored subsystems in which the target set X' can be reached limit-surely, i.e., let $J = \{i : 0 \leq i \leq m - 1 \mid \mathbb{P}^+(M_{S'}, r_i, \llbracket \Diamond X' \rrbracket) = 1\}$ be the set of good colors and let $R = \{r_j \mid j \in J\}$, and $\bar{R} = \{r_j \mid j \notin J\}$.

Further, let X'_f be the restriction of X' to the finite uncolored part of S' (i.e., only those parts added in line 25 of Alg. 1).

We now establish the existence of certain strategies in subsystems of S' . These will later serve as building blocks for our strategies on M_S .

Since we assumed that $\mathbb{P}^+(M_{S'}, c'_0, \llbracket \Diamond X' \rrbracket) = 1$, there exists a family of strategies that makes the probability of reaching X' arbitrarily close to one. In particular, they must also make the probability of reaching configurations in \bar{R} arbitrarily close to zero. Thus we obtain $\mathbb{P}^+(M_{S'}, c'_0, \llbracket \Diamond X'_f \cup R \rrbracket) = 1$, i.e., we can limit-surely reach $X'_f \cup R$. Since, for this objective, only the finite uncolored part of $M_{S'}$ is relevant, this is a problem for a finite-state MDP and limit-surely and almost-surely coincide. So there exists a partial strategy σ , for the uncol-

ored part of $M_{S'}$, such that, starting in c'_0 , we almost-surely reach $X'_f \cup R$, i.e., $\mathbb{P}(M_{S'}, c'_0, \sigma, \llbracket \Diamond X'_f \cup R \rrbracket) = 1$.

In each of the good colored subsystems we can limit-surely reach X' , i.e., for every $r_j \in R$ we have $\mathbb{P}^+(M_{S'}, r_i, \llbracket \Diamond X' \rrbracket) = 1$. So for every $\epsilon > 0$ there exists a strategy σ_i^ϵ such that $\mathbb{P}(M_{S'}, r_i, \sigma_i^\epsilon, \llbracket \Diamond X' \rrbracket) \geq 1 - \epsilon$. Consider the computation tree of the game on $M_{S'}$ from r_i when playing according to σ_i^ϵ and its restriction to some finite depth d . Let $\mathbb{P}_d(M_{S'}, r_i, \sigma_i^\epsilon, \llbracket \Diamond X' \rrbracket)$ be the probability that the objective is reached already during the first d steps of the game. We have $\mathbb{P}_d(M_{S'}, r_i, \sigma_i^\epsilon, \llbracket \Diamond X' \rrbracket) \leq \mathbb{P}(M_{S'}, r_i, \sigma_i^\epsilon, \llbracket \Diamond X' \rrbracket)$, but $\lim_{d \rightarrow \infty} \mathbb{P}_d(M_{S'}, r_i, \sigma_i^\epsilon, \llbracket \Diamond X' \rrbracket) = \mathbb{P}(M_{S'}, r_i, \sigma_i^\epsilon, \llbracket \Diamond X' \rrbracket) \geq 1 - \epsilon$. Thus for every color $i \in J$ and every $\epsilon > 0$ there exists a number $d(i, \epsilon)$ s.t. $\mathbb{P}_{d(i, \epsilon)}(M_{S'}, r_i, \sigma_i^\epsilon, \llbracket \Diamond X' \rrbracket) \geq 1 - 2\epsilon$.

Since configurations of $M_{S'}$ are obtained by projecting configurations of M_S , we can go the reverse direction by replacing the missing component in an $M_{S'}$ configuration by a given number. Given an $M_{S'}$ -configuration r_i and a number $d(i, \epsilon)$ we obtain an M_S -configuration $s_i(d(i, \epsilon))$ by replacing the missing $k(i)$ -th component of r_i by $d(i, \epsilon)$. Let $\alpha \in \mathbb{N}$ be the maximal constant appearing in any transition in S , i.e., the maximal possible change in any counter in a single step. Since a single step in M_S can only change a counter by $\leq \alpha$, the $k(i)$ -th component of $s_i(\alpha * d(i, \epsilon))$ cannot be exhausted during the first $d(i, \epsilon)$ steps of the game on M_S starting at $s_i(\alpha * d(i, \epsilon))$. Thus we can use the same strategy σ_i^ϵ in the game from $s_i(\alpha * d(i, \epsilon))$ on M_S and obtain $\mathbb{P}(M_S, s_i(\alpha * d(i, \epsilon)), \sigma_i^\epsilon, \llbracket \Diamond X \rrbracket) \geq 1 - 2\epsilon$. Intuitively, the number $\alpha * d(i, \epsilon)$ is big enough to allow playing the game for sufficiently many steps to make the probability of success close to 1.

Using the strategy σ above and the strategies σ_i^ϵ , we now define a new family of strategies σ_ϵ for every $\epsilon > 0$ for the game on M_S from c_0 . Given $\epsilon > 0$, we let $d(\epsilon) = \alpha * \max_{i \in J} d(i, \epsilon)$ (a number that is big enough for each projected component).

Playing from c_0 in M_S , the strategy σ_ϵ behaves as follows. First it plays like strategy σ in the corresponding game from c'_0 on $M_{S'}$. (Function λ connects the corresponding configurations in the two games.) When the game in $M_{S'}$ reaches a configuration r_i then there are two cases: If the configuration in M_S is $\geq s_i(d(\epsilon))$ then σ_ϵ henceforth plays like σ_i^ϵ , which ensures to reach the target X with probability $\geq 1 - 2\epsilon$. Otherwise, the configuration in M_S is still too small to switch to σ_i^ϵ . In this case, σ_ϵ continues to play like σ plays from the previously visited smaller configuration in the uncolored part of $M_{S'}$ (see line 7 of the algorithm). This is possible, because the game is monotone and larger configurations always benefit Player 1. So the game on M_S continues with a configuration that is larger (at least on component $k(i)$) than the corresponding game on $M_{S'}$, i.e., component $k(i)$ is pumped. Since we know that σ on $M_{S'}$ will almost surely visit X'_f or R , we obtain that σ_ϵ on M_S will almost surely eventually visit X or some configuration $\geq s_i(d(\epsilon))$ for $i \in J$ (and from there achieve to reach the target with probability $\geq 1 - 2\epsilon$). Since every weighted average of probabilities $\geq 1 - 2\epsilon$ is still $\geq 1 - 2\epsilon$, we obtain $\mathbb{P}(M_S, c_0, \sigma_\epsilon, \llbracket \Diamond X \rrbracket) \geq 1 - 2\epsilon$ and thus $\mathbb{P}^+(M_S, c_0, \llbracket \Diamond X \rrbracket) = 1$. \square

B.9 Proof of Theorem 9

Proof. Let $S = \langle Q, Q_1, Q_P, T, \tau \rangle$ be 1-VASS-MDP of dimension $d > 0$ with initial configuration $c_0 = \langle q_0, \mathbf{v} \rangle$ and $X \subseteq Q$ a set of target states. We show decidability of $\mathbb{P}^+(M_S, c_0, \llbracket \Diamond X \rrbracket) = 1$ by induction on d .

Base case $d = 0$. If S has 0 counters then M_S is a finite-state MDP and thus limit sure reachability coincides with almost sure reachability, which is decidable.

Inductive step. We apply Algorithm 1, which terminates by Lemma 6, and obtain a new instance of the 1-VASS-MDP limit sure reachability problem of dimension $d - 1$: $S' = \langle Q', Q'_1, Q'_P, T', \tau' \rangle$ with initial configuration $c'_0 = \langle q'_0, \mathbf{0} \rangle$ and set of target states $X' \subseteq Q'$. By Lemma 16 and Lemma 17, we have $\mathbb{P}^+(M_S, c_0, \llbracket \Diamond X \rrbracket) = 1 \Leftrightarrow \mathbb{P}^+(M_{S'}, c'_0, \llbracket \Diamond X' \rrbracket) = 1$. By induction hypothesis, $\mathbb{P}^+(M_{S'}, c'_0, \llbracket \Diamond X' \rrbracket) = 1$ is decidable and the result follows. \square